Introduction to AI and ML

Fredrik Heintz
Dept. of Computer Science, Linköping University
Dept. of Communication and Media, Lund University
fredrik.heintz@liu.se / fredrik.heintz@kom.lu.se
@FredrikHeintz

LINKÖPINGS UNIVERSITET

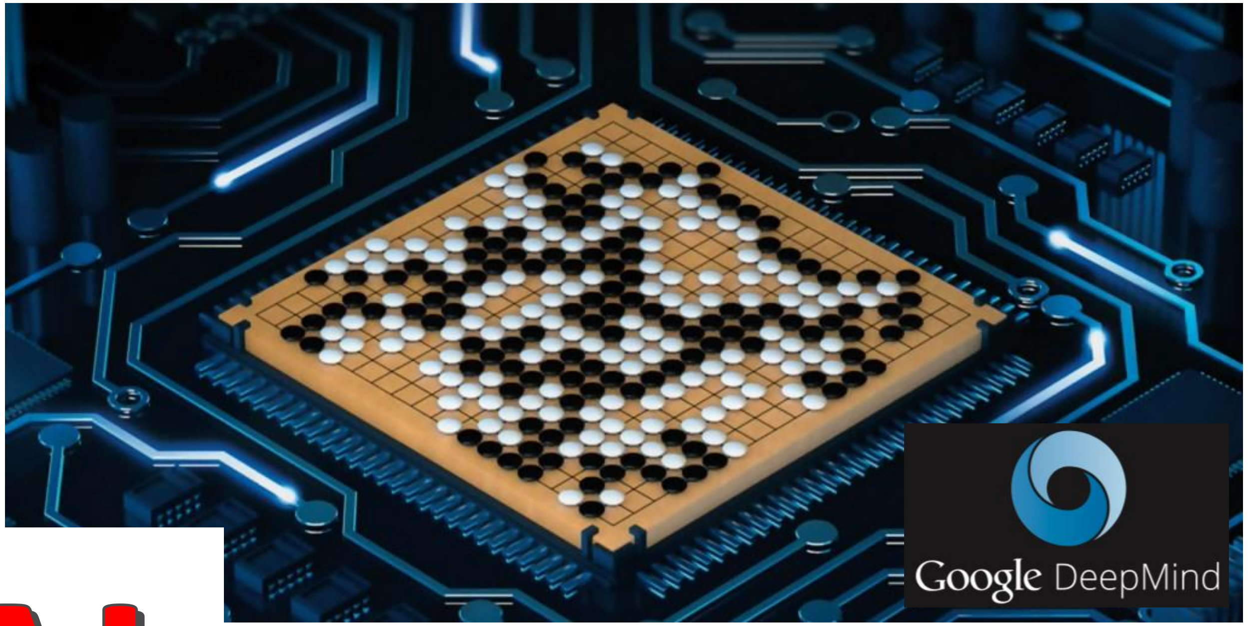
Google DeepMind

AI


THINK
$300,000
$1,000,000
$200,000
KEN
WATSON
BRAD


BostonDynamics
Boston Dynamics

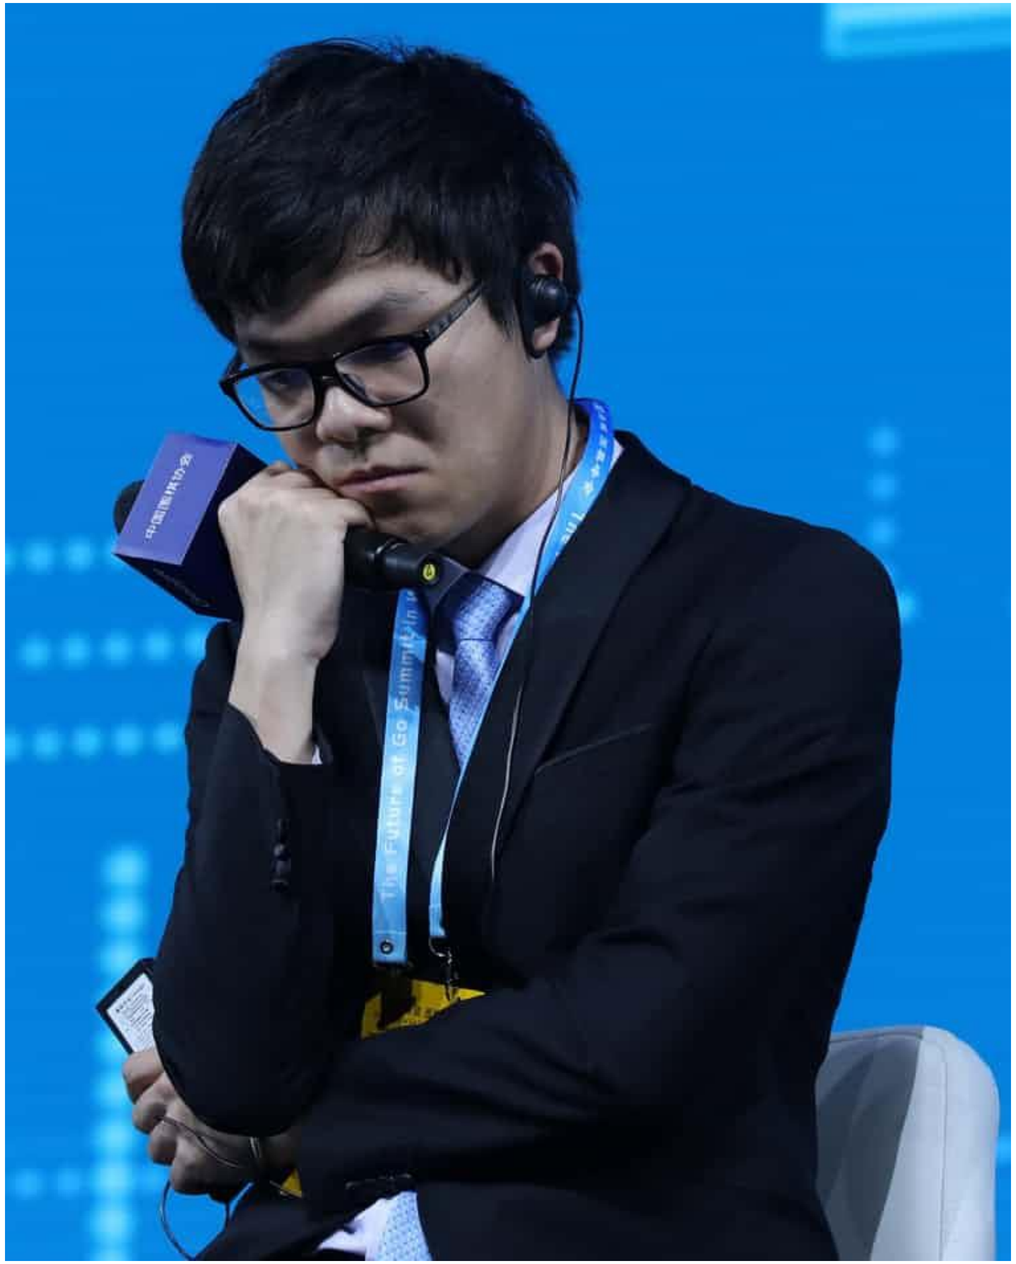


柯洁 KE JIE
00:46:57

AlphaGo Overview

based on: Silver, D. et al. Nature Vol 529, 2016
copyright: Bob van den Hoek, 2016

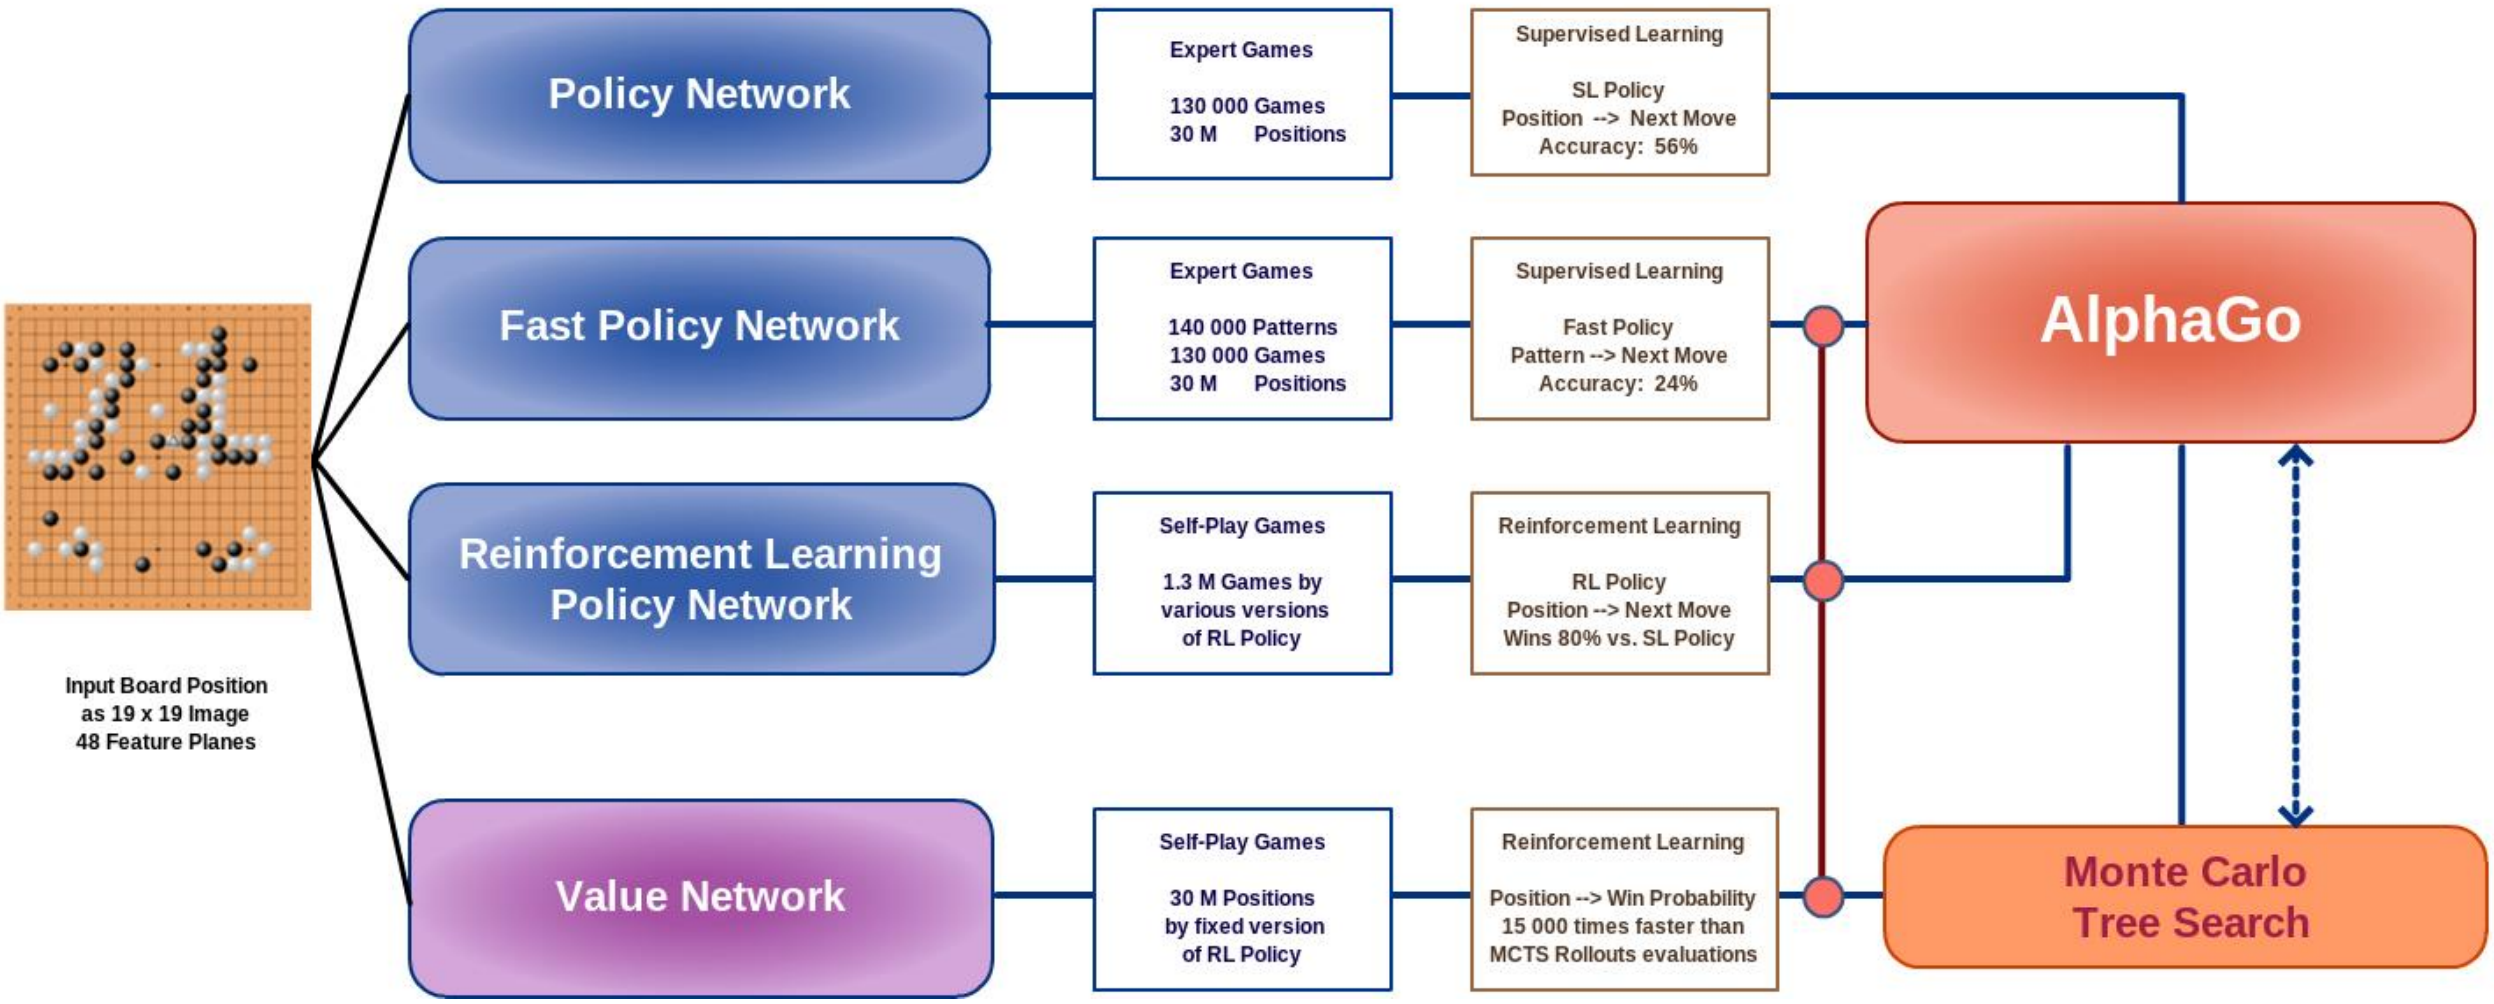

http://deeplearningskysthelimit.blogspot.se/2016/04/part-2-alphago-under-magnifying-glass.html

Artificial Intelligence – What is it? – Definitions

"Artificial Intelligence is the **science** and **engineering** of **making intelligent machines**, especially intelligent computer programs."
- John McCarthy, Stanford

"Artificial intelligence (AI) refers to **systems that display intelligent behaviour** by analysing their environment and taking actions – with some degree of **autonomy** – to achieve specific **goals**."
- EU Communication 25 April 2018

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

Artificial Intelligence – Four Views

	Empirical Sciences Fidelity to human performance **Human-Centered**	Mathematics/Engineering Ideal concept of Intelligence **Rationality-Centered**
Thought Processes Reasoning	Systems that think like humans	Systems that think rationally
	”The exciting new effort to make computers think. . .machines with minds, in the full and literal sense.” (Haugeland, 1985) ”[The automation of] activities that we associate with human thinking, activities such as decision-making, problem solving, learning...”(Bellman, 1978)	”The study of mental faculties through the use of computational models.” (Charniak and McDermott, 1985) ”The study of computations that make it possible to perceive, reason, and act.” (Winston, 1992)
Behavior	Systems that act like humans	Systems that act rationally
	”The art of creating machines that perform functions that require intelligence when performed by people.” (Kurzweil, 1990) ”The study of how to make computers do things at which, at the moment, people are better.” (Rich and Knight, 1991)	”Computational Intelligence is the study of the design of intelligent agents.” (Poole et al., 1998) ”AI . . . Is concerned with intelligent behavior in artifacts.” (Nilsson, 1998)

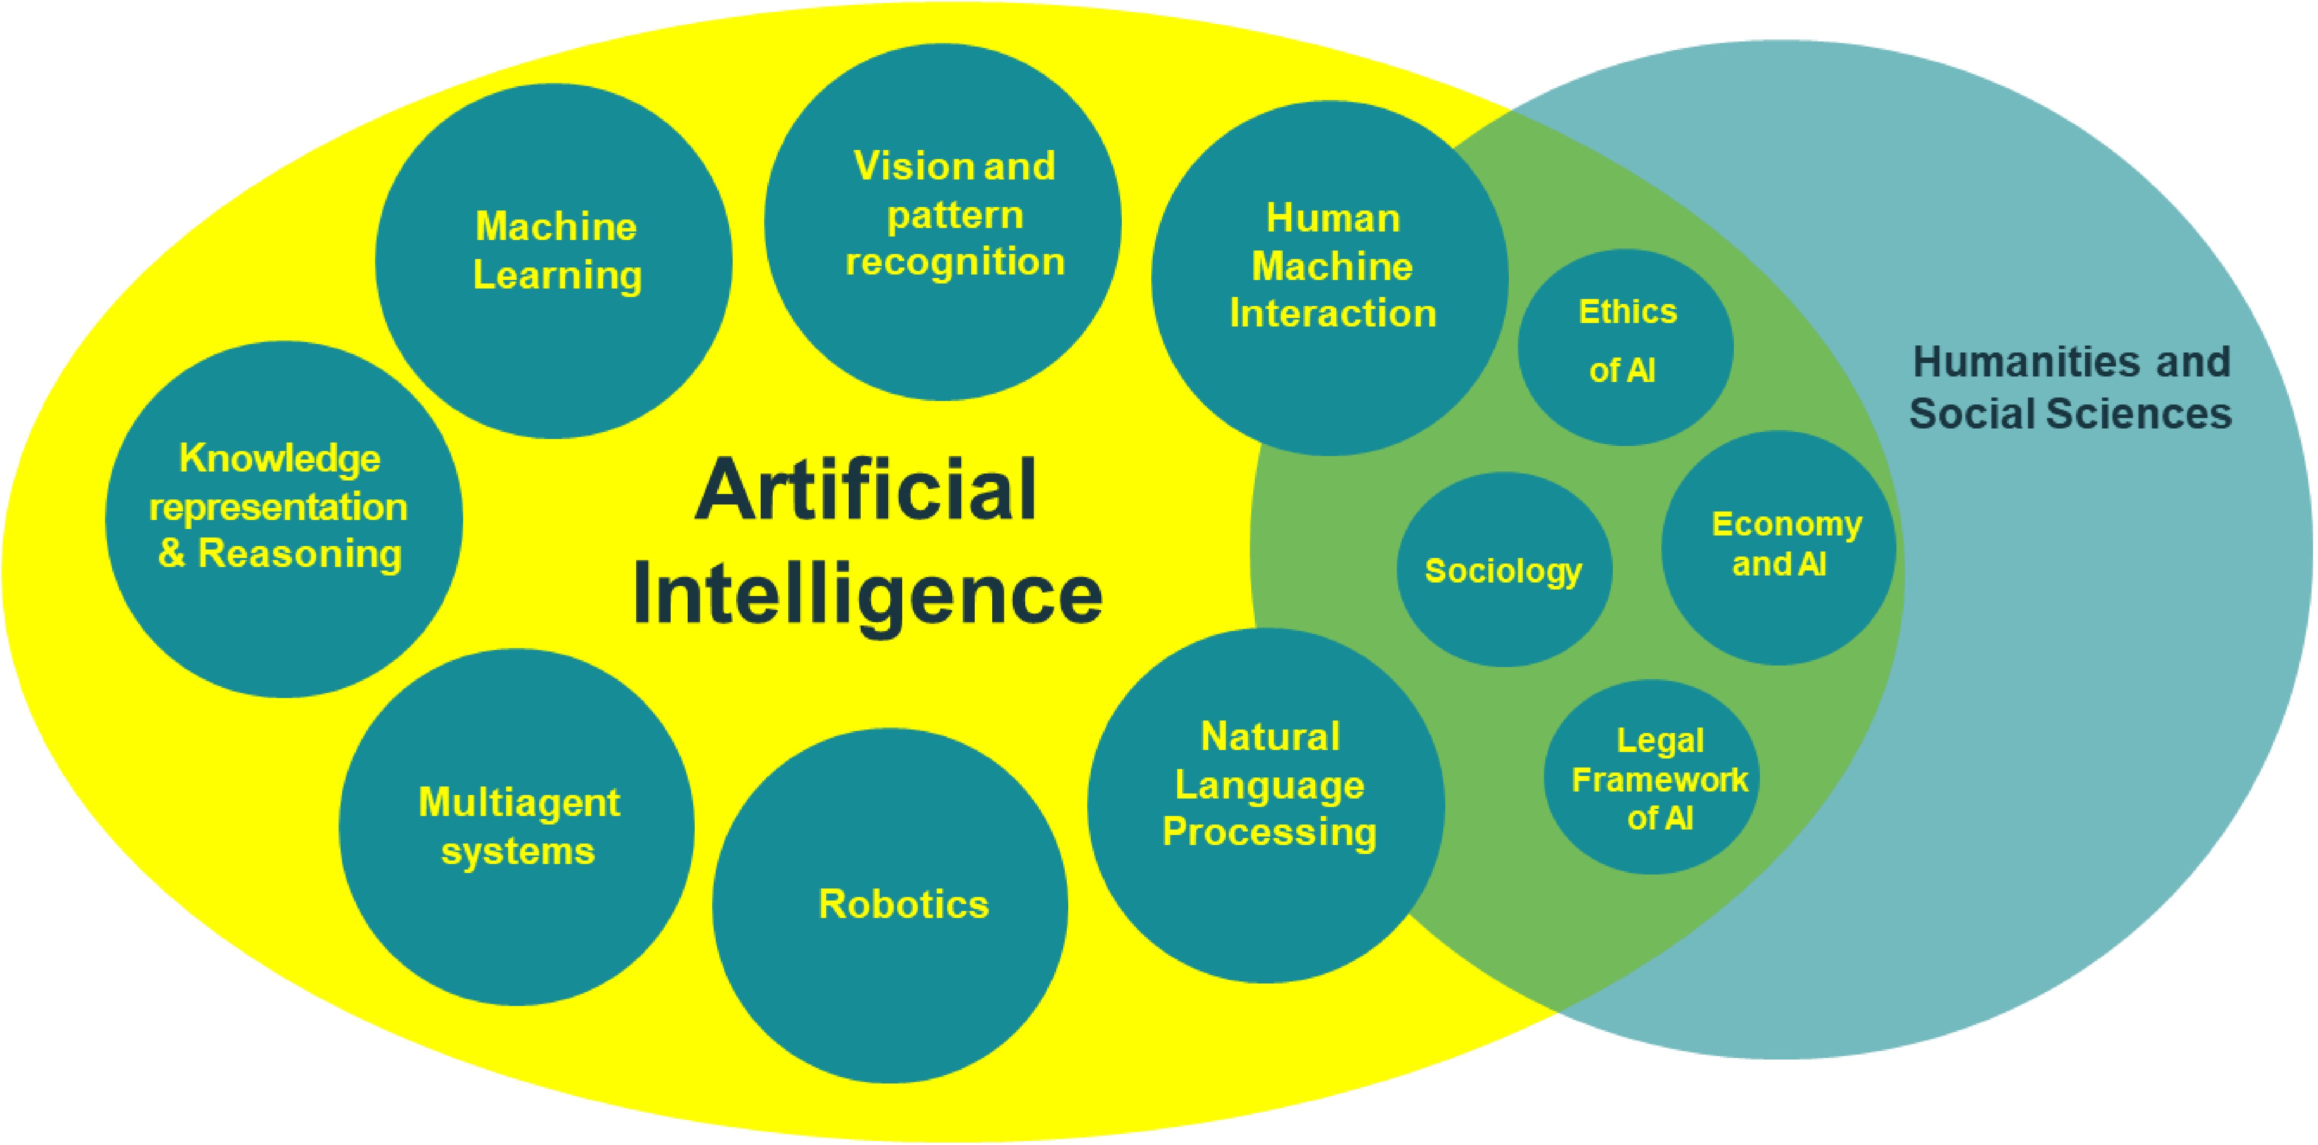
Machine Learning
Vision and pattern recognition
Human Machine Interaction
Knowledge representation & Reasoning
Artificial Intelligence
Ethics of AI
Humanities and Social Sciences
Sociology
Economy and AI
Multiagent systems
Robotics
Natural Language Processing
Legal Framework of AI

Russell & Norvig – Artificial Intelligence: A Modern Approach

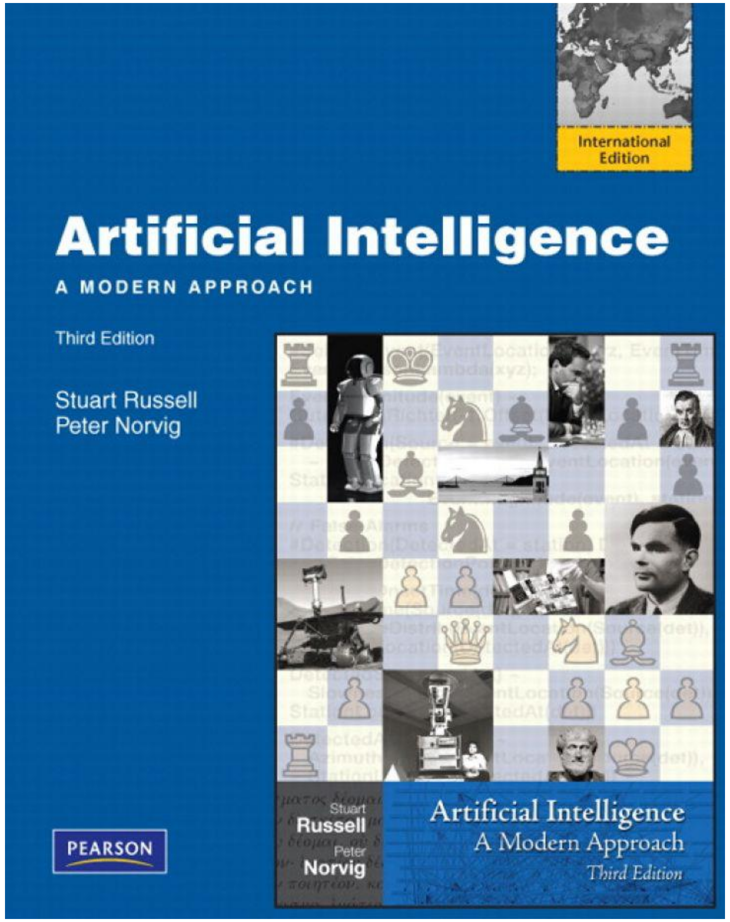

http://aima.cs.berkeley.edu/

AI as Intelligent Agents

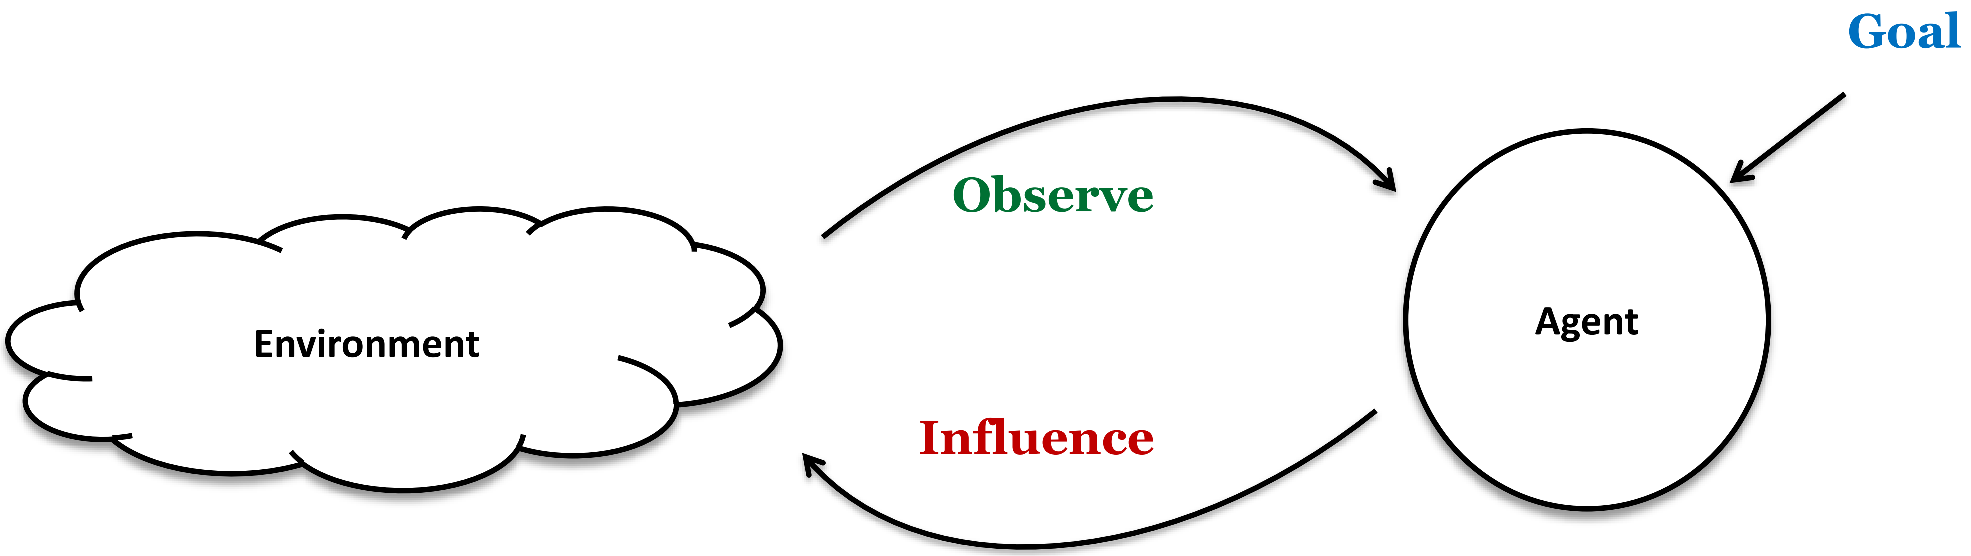

AI as Intelligent Agents

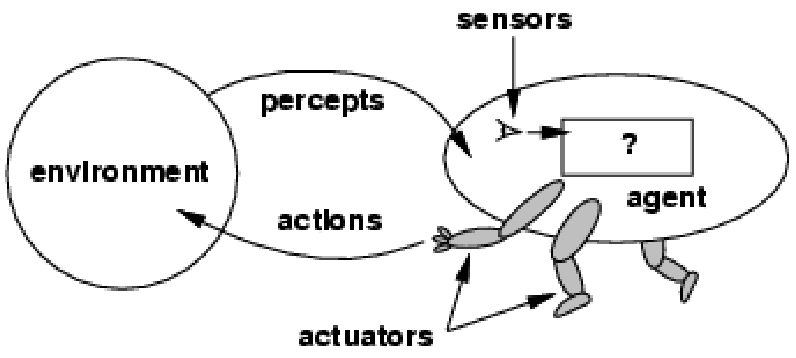

- The agent function maps from percept histories to actions:

$$[f: \mathcal{P}^{\star} \rightarrow \mathcal{A}]$$

- The agent program runs on the physical architecture to produce f
- agent = architecture + program

AI as Intelligent Agents

1. While true {
2. **Sense** the world– (a) sensors, (b) communication, (c) supervisor input
    1. Form perceptions– (a) concept triggering, (b) propioception
    2. update beliefs (belief revision)
    3. update internal world model– (a) map, (b) localization, (c) relationships and attributes
3. **Think** about options, desires, intentions, and actions
    1. Revise desirable options and select one
    2. Deliberate about what intention to achieve next;
    3. Revise and update plan
    4. use means-ends reasoning to get a plan for the intention;
- 4. **Act**
    1. Revise intentions and select an intention to manifest
    2. execute the plan
    3. Suppress less important behaviors
    4. Start control of actuators
- 5. **Pause**
    1. until the world changes
    2. Communicate
    3. Generate and deliver user feedback

    }

The-
Frame problem
Action selection problem

Replanning problem
Envisionment problem

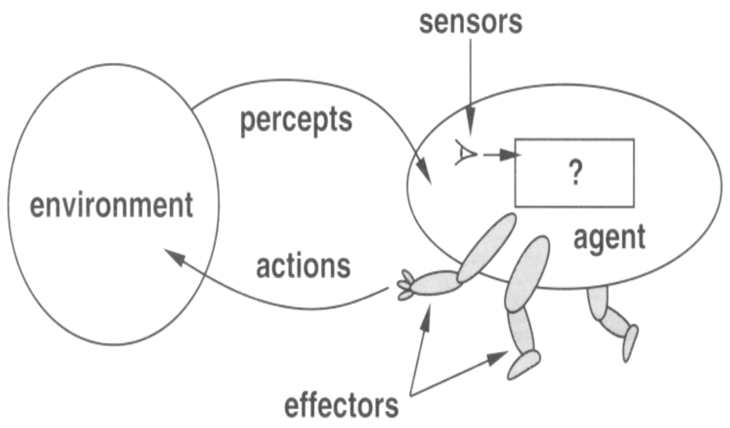

HDRC3: A Distributed Hybrid Deliberative/Reactive Architecture for Autonomous Systems

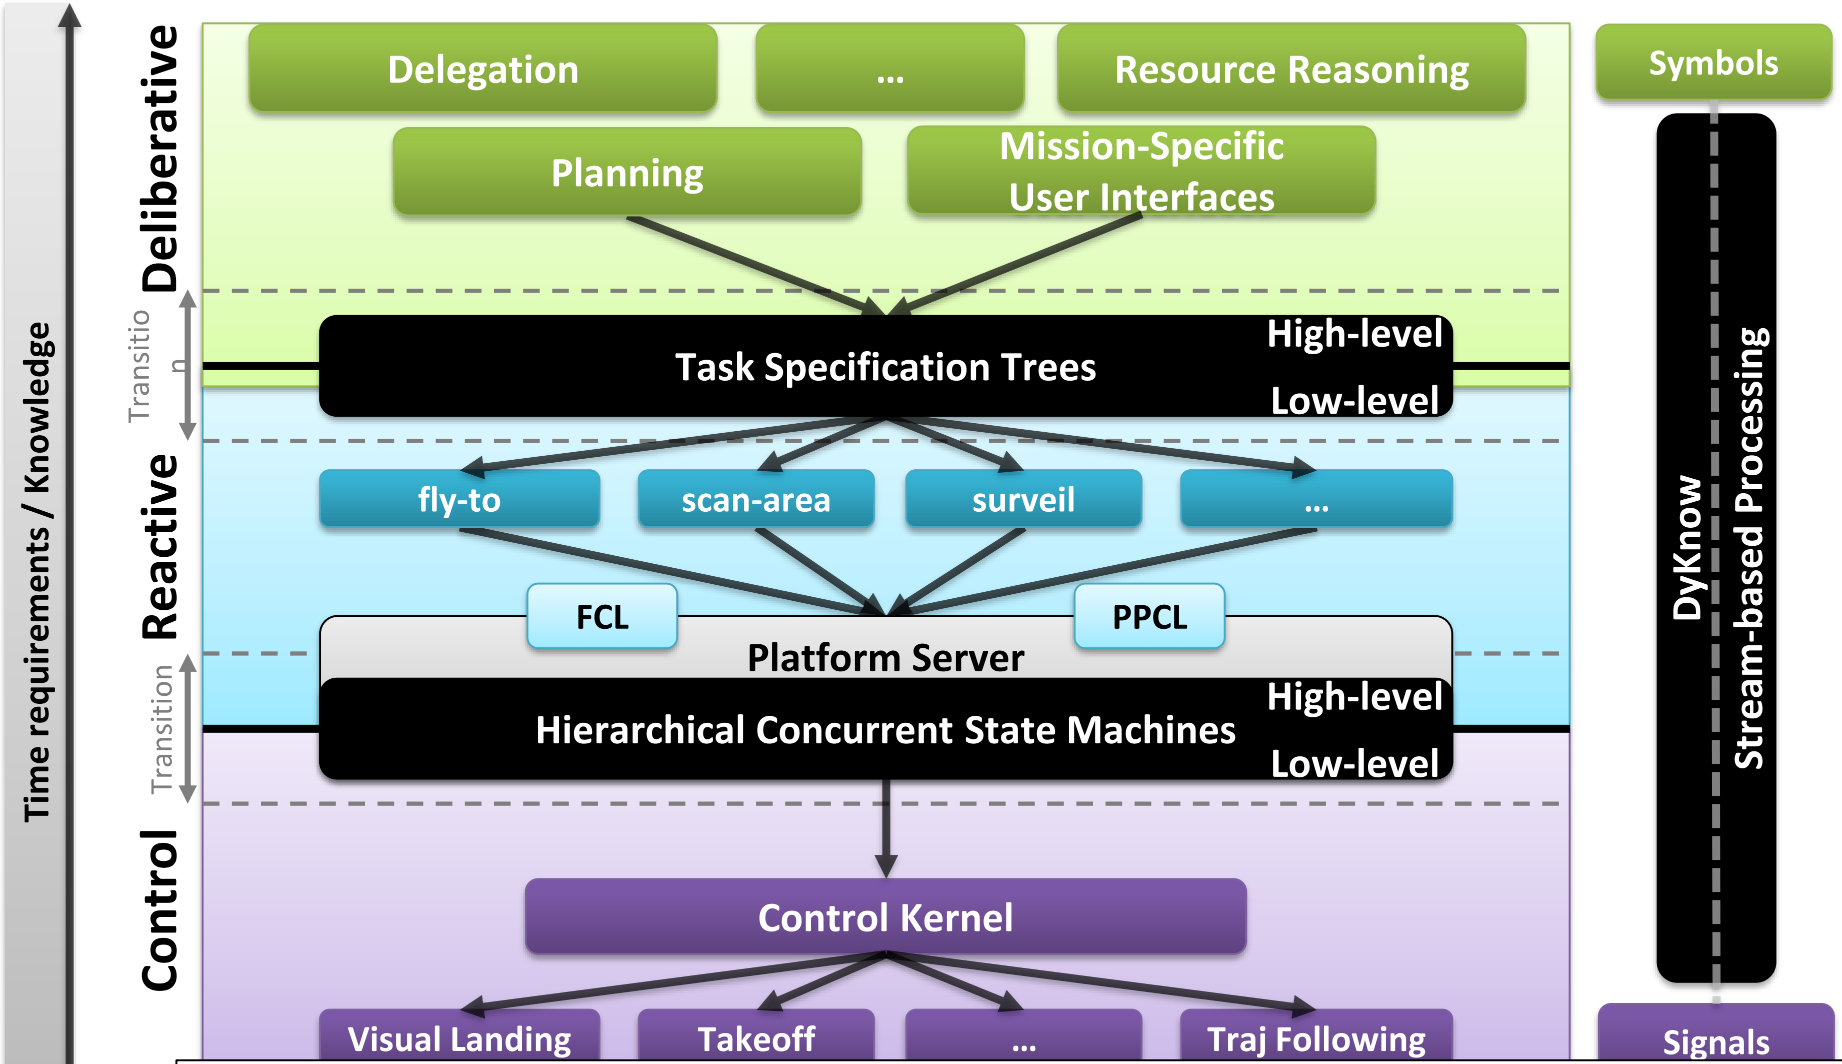

P. Doherty, J. Kvarnström, M. Wzorek, P. Rudol, F. Heintz and G. Conte. 2014.
HDRC3 - A Distributed Hybrid Deliberative/Reactive Architecture for Unmanned Aircraft Systems.
In K. Valavanis, G. Vachtsevanos, editors, Handbook of Unmanned Aerial Vehicles, pages 849–952.

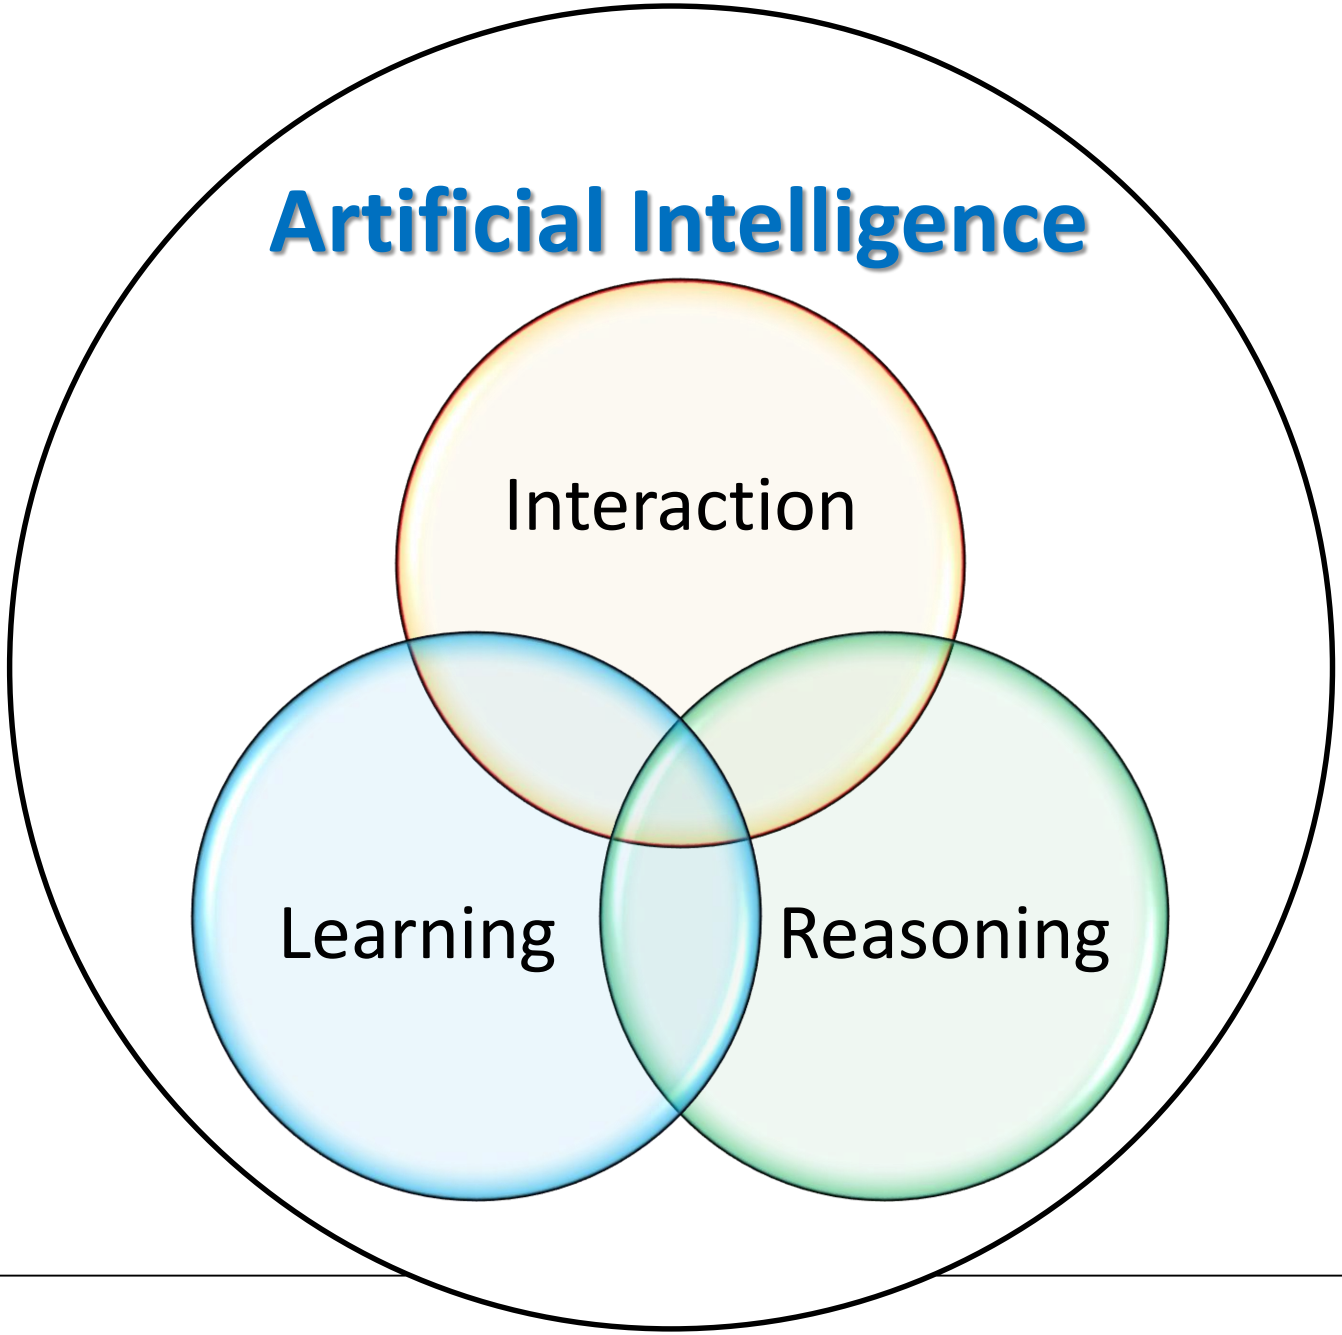
Artificial Intelligence
Interaction
Learning
Reasoning

Erik Sandewall

https://liu.se/artikel/erik-sandewall---den-envise-visionaren

Moore's Law – The number of transistors on integrated circuit chips (1971-2018)

Our World in Data

Moore's law describes the empirical regularity that the number of transistors on integrated circuits doubles approximately every two years. This advancement is important as other aspects of technological progress – such as processing speed or the price of electronic products – are linked to Moore's law.

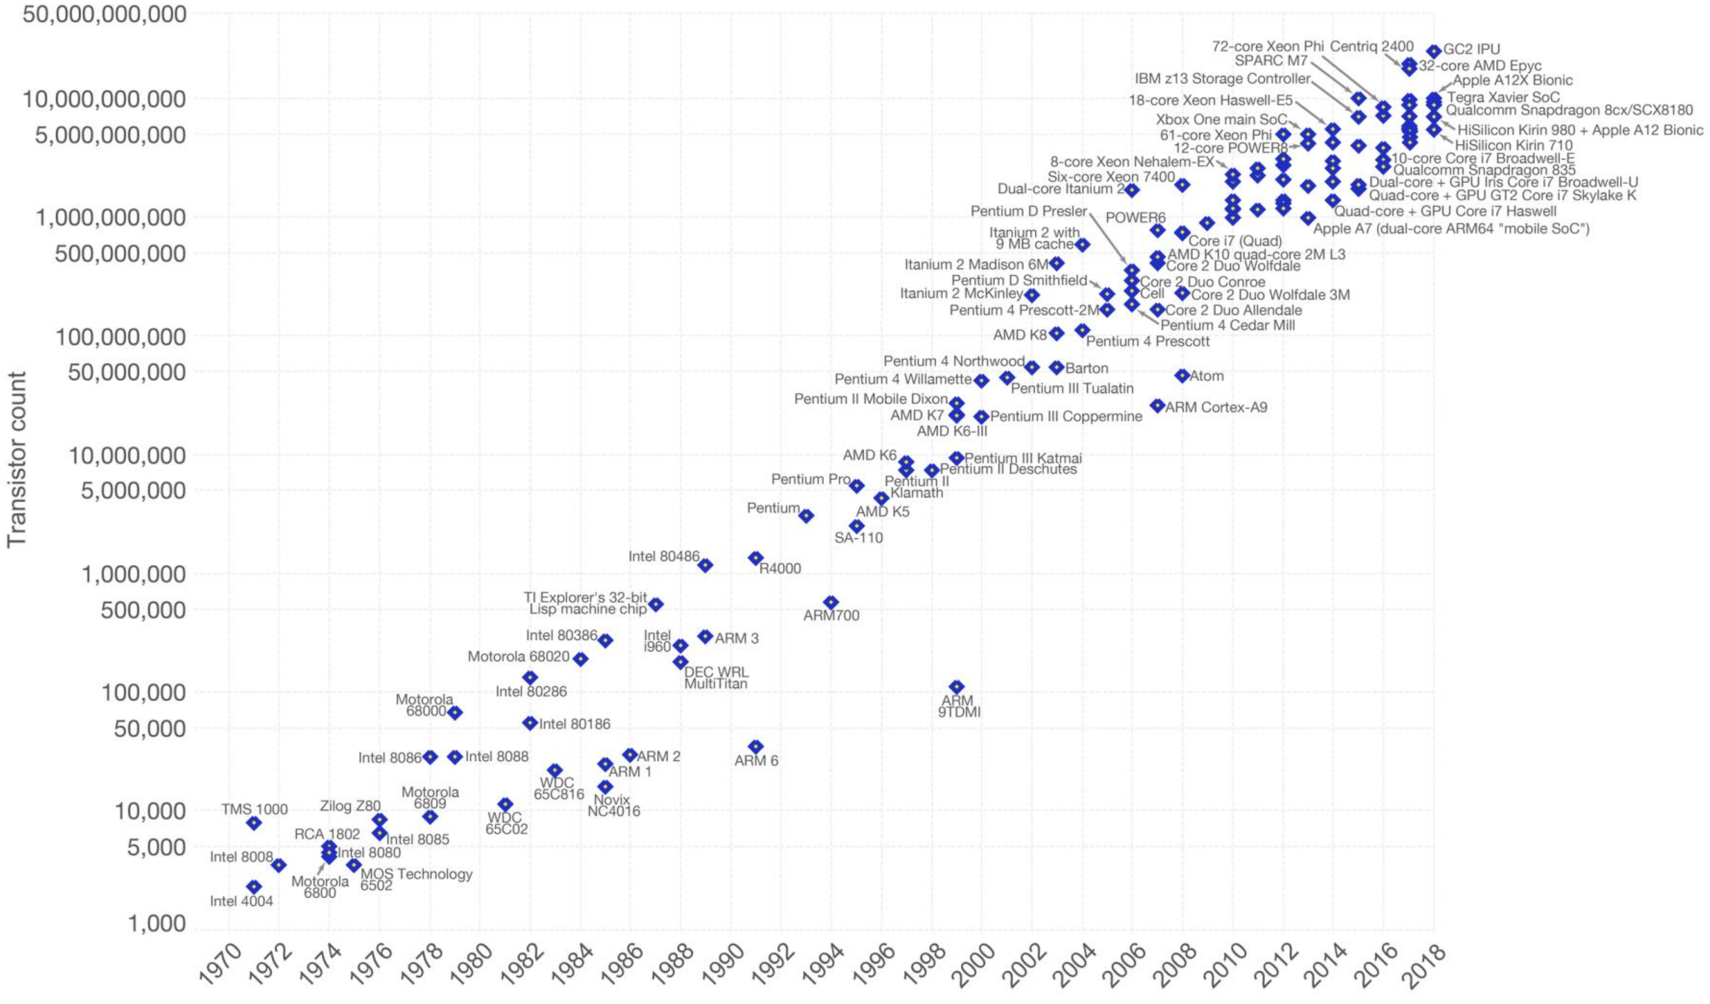

Data source: Wikipedia (https://en.wikipedia.org/wiki/Transistor_count)
The data visualization is available at OurWorldinData.org. There you find more visualizations and research on this topic.

https://en.wikipedia.org/wiki/Moore's_law


You
Tube
Bē
Meetup

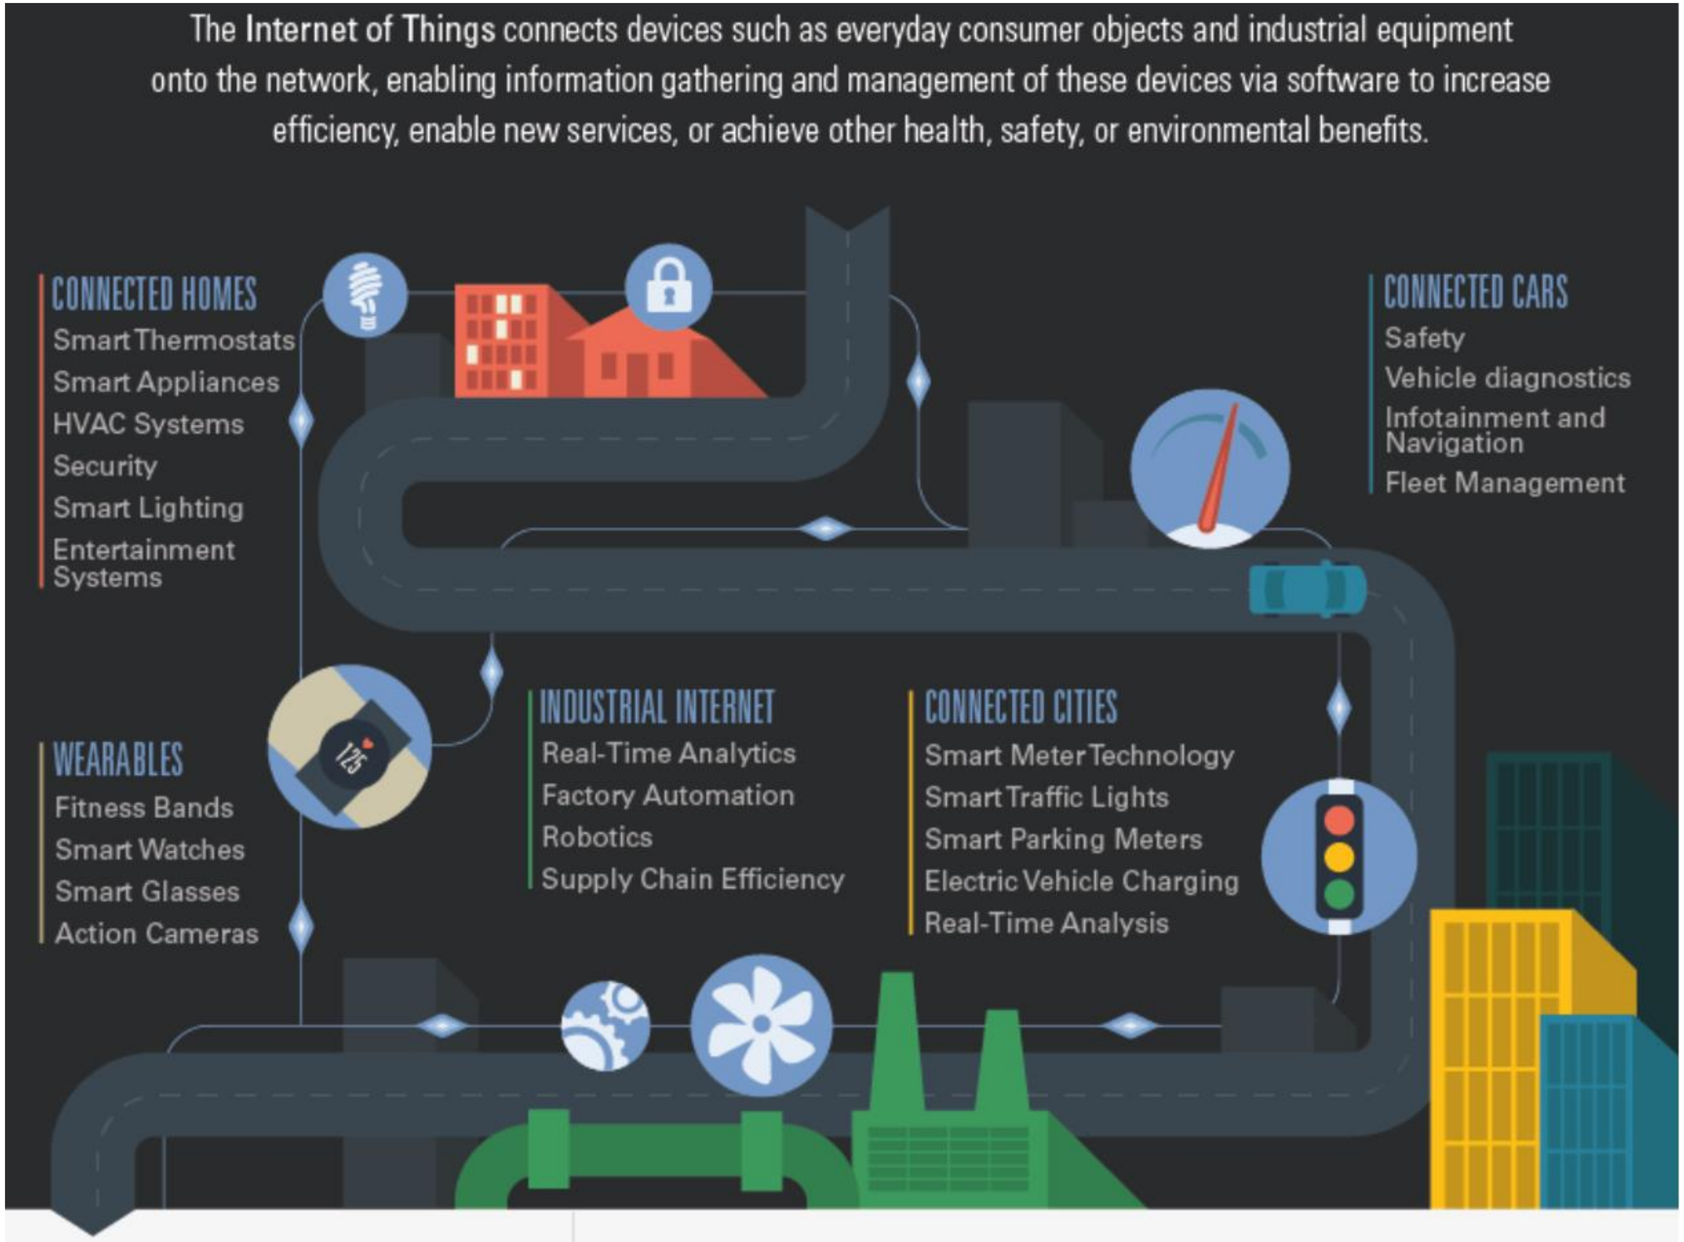
The Internet of Things connects devices such as everyday consumer objects and industrial equipment onto the network, enabling information gathering and management of these devices via software to increase efficiency, enable new services, or achieve other health, safety, or environmental benefits.
CONNECTED HOMES
Smart Thermostats
Smart Appliances
HVAC Systems
Security
Smart Lighting
Entertainment Systems
CONNECTED CARS
Safety
Vehicle diagnostics
Infotainment and Navigation
Fleet Management
WEARABLES
Fitness Bands
Smart Watches
Smart Glasses
Action Cameras
INDUSTRIAL INTERNET
Real-Time Analytics
Factory Automation
Robotics
Supply Chain Efficiency
CONNECTED CITIES
Smart Meter Technology
Smart Traffic Lights
Smart Parking Meters
Electric Vehicle Charging
Real-Time Analysis

The growth of human and machine-generated data

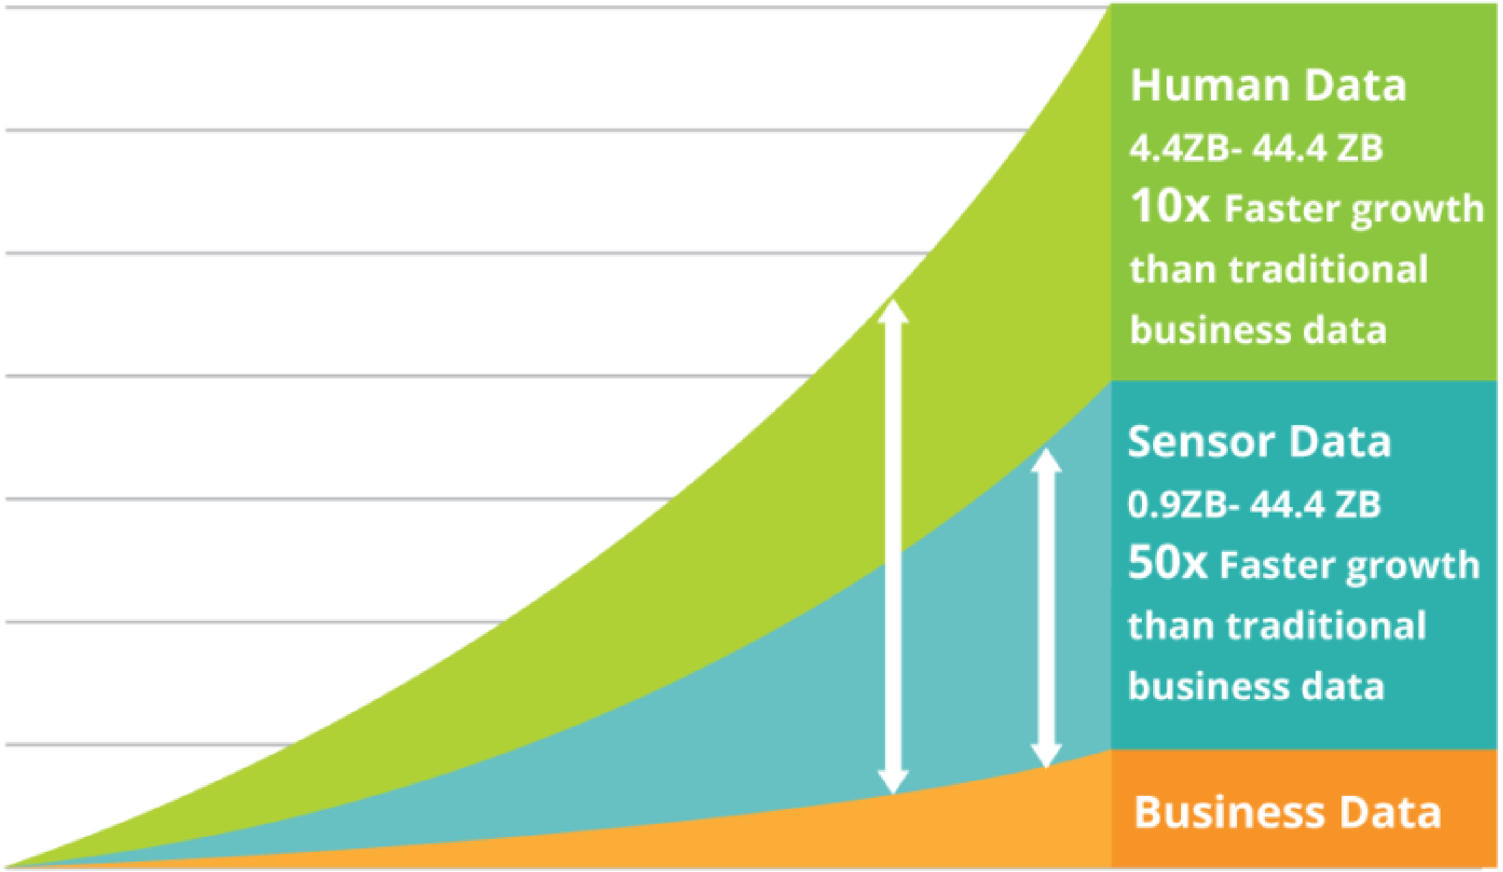

Source: Inside big data

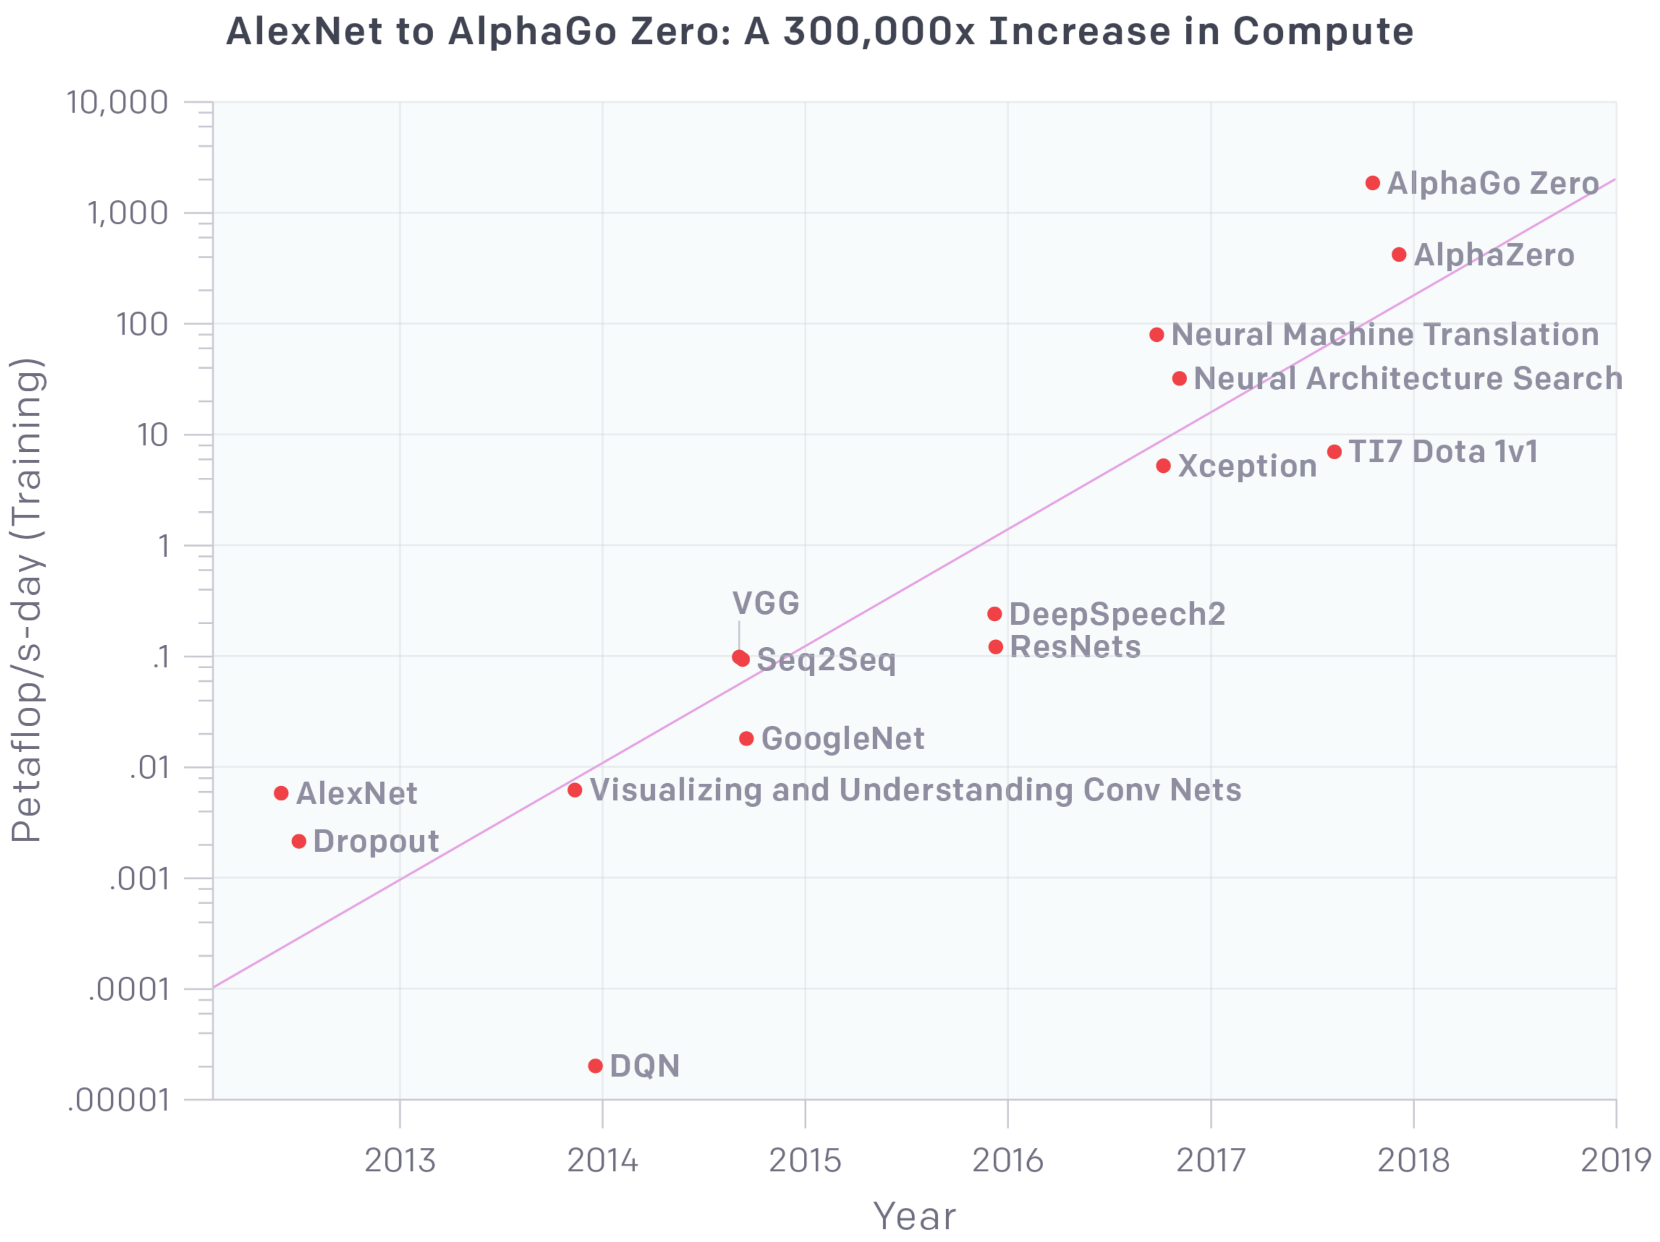

LINKÖPINGS UNIVERSITET

https://openai.com/blog/ai-and-compute/

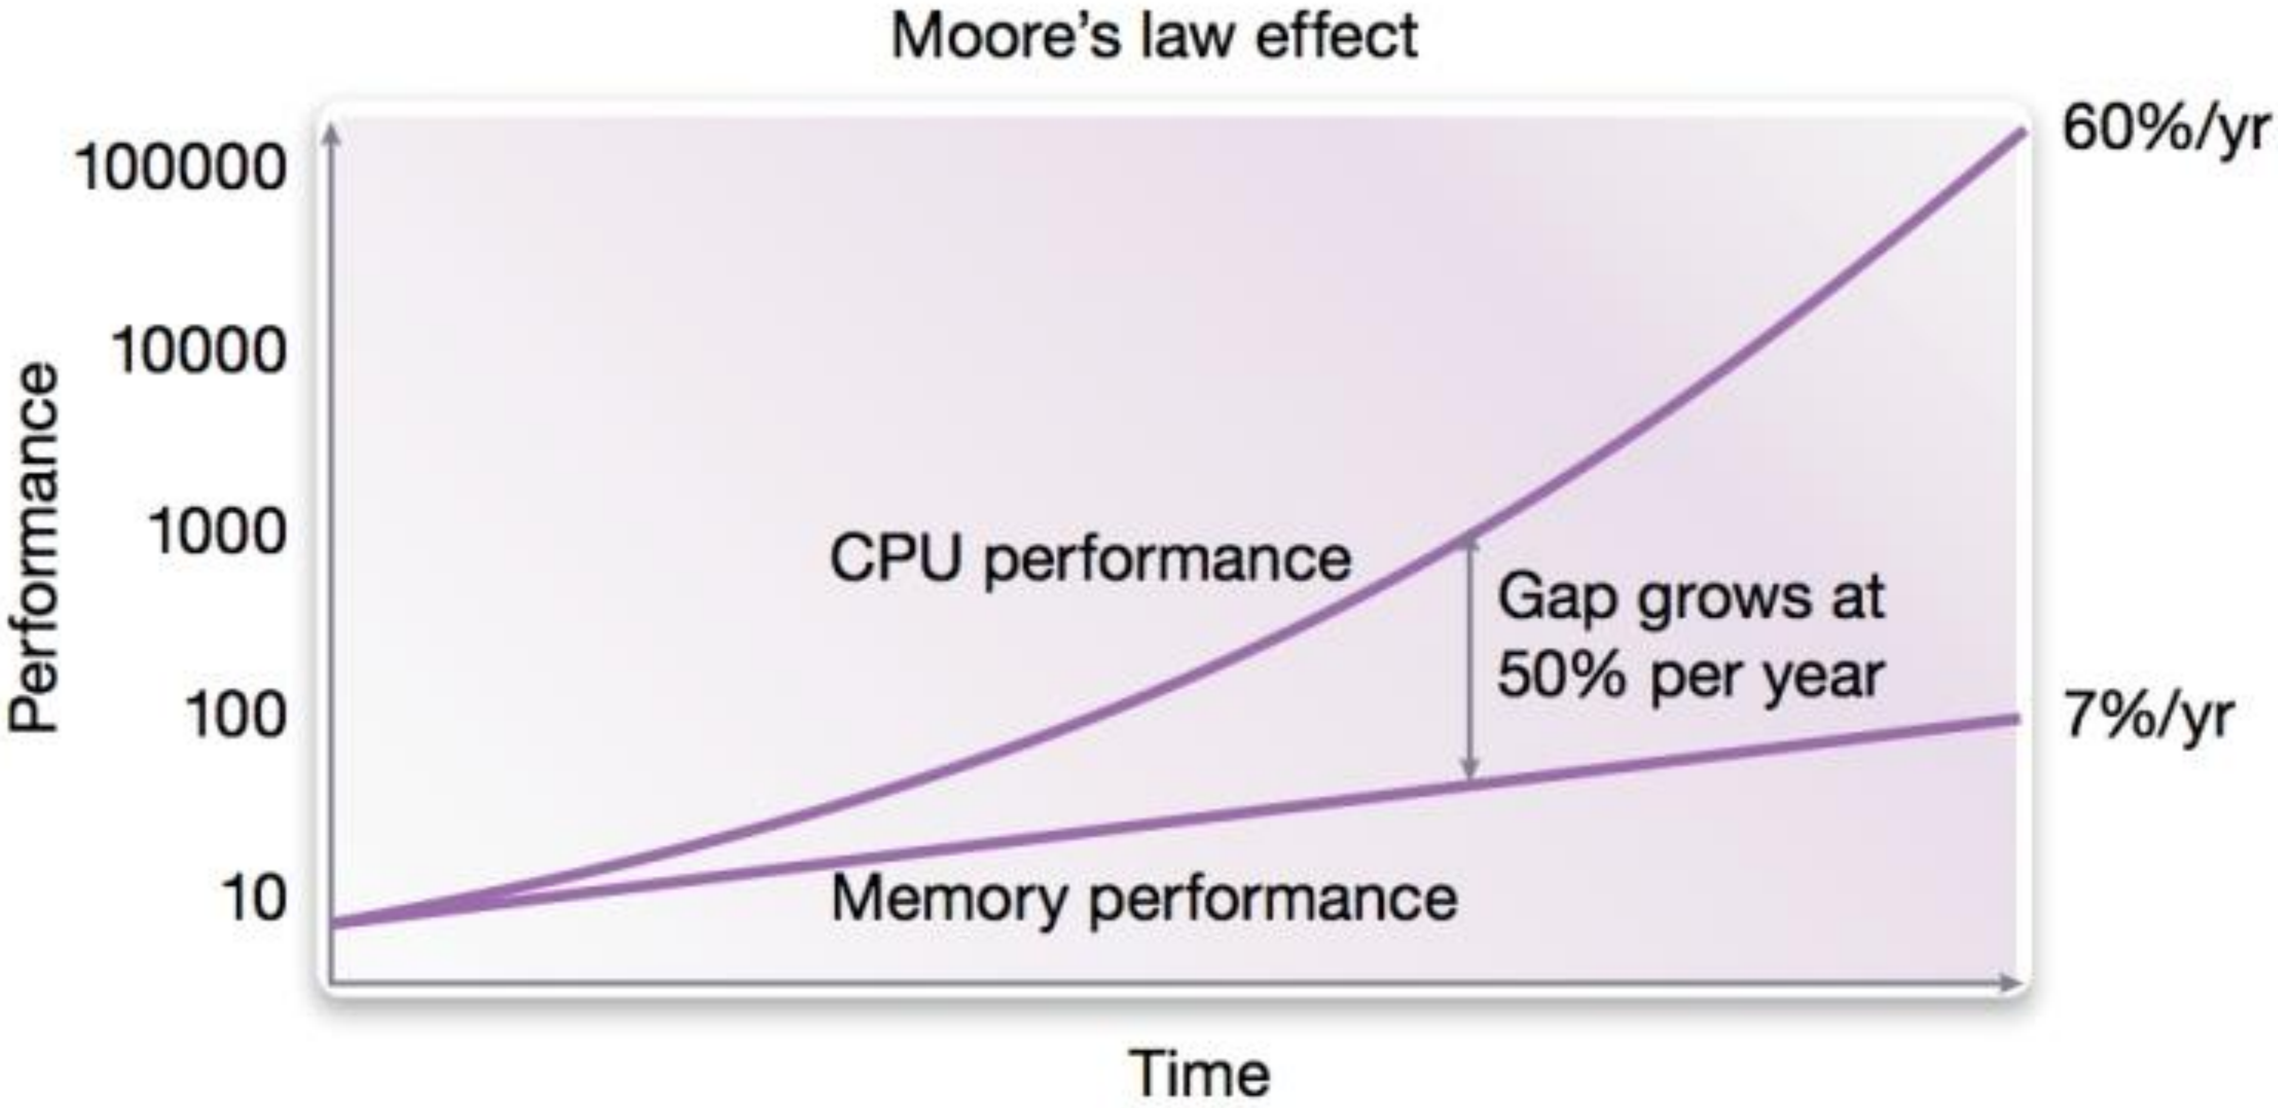

https://www.extremetech.com/computing/233691-phase-change-memory-can-operate-thousands-of-times-faster-than-current-ram

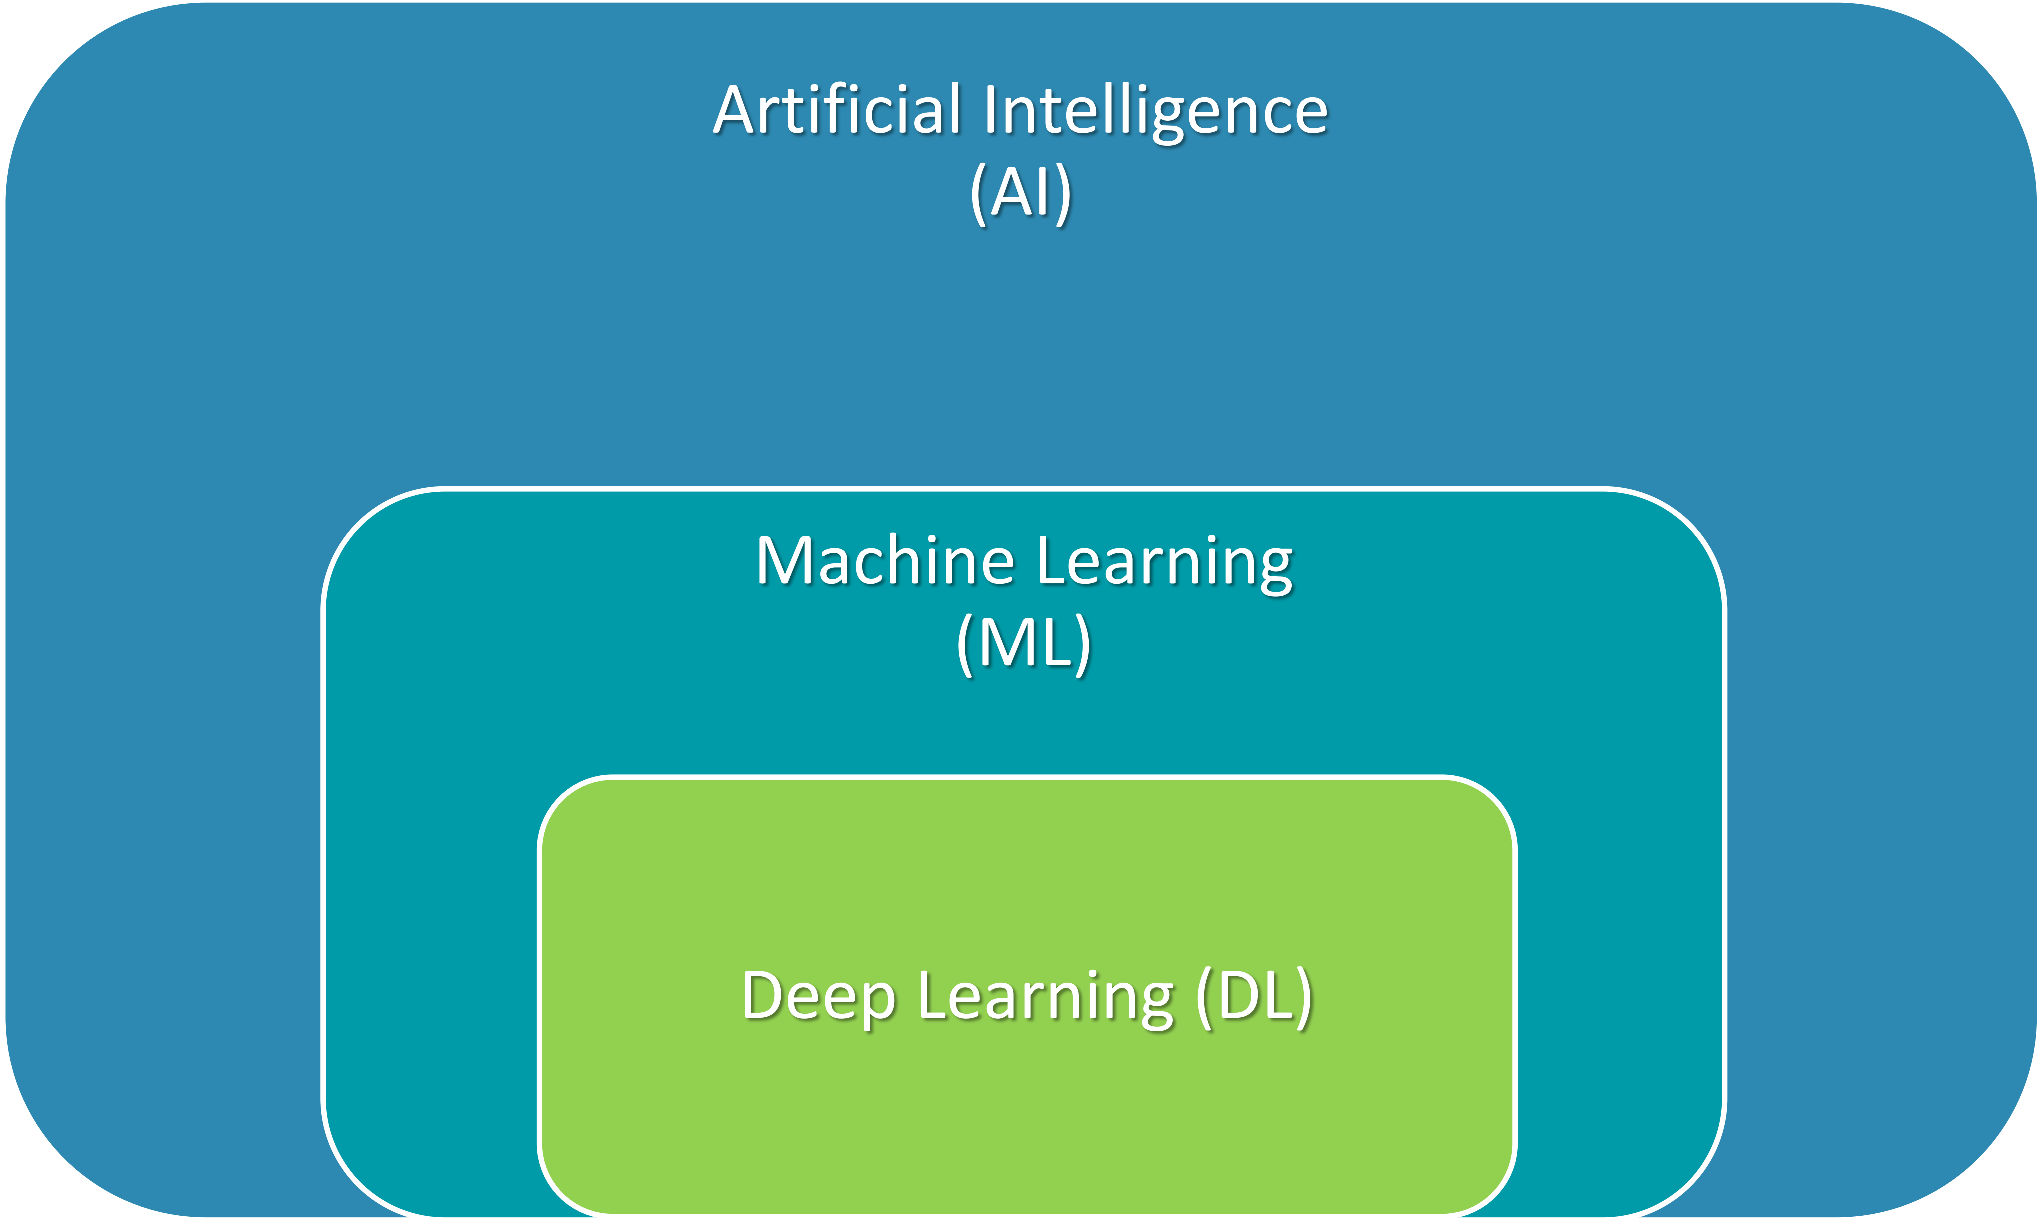
Artificial Intelligence
(AI)
Machine Learning
(ML)
Deep Learning (DL)

Algorithmic, Knowledge-Based and Learning-Based AI

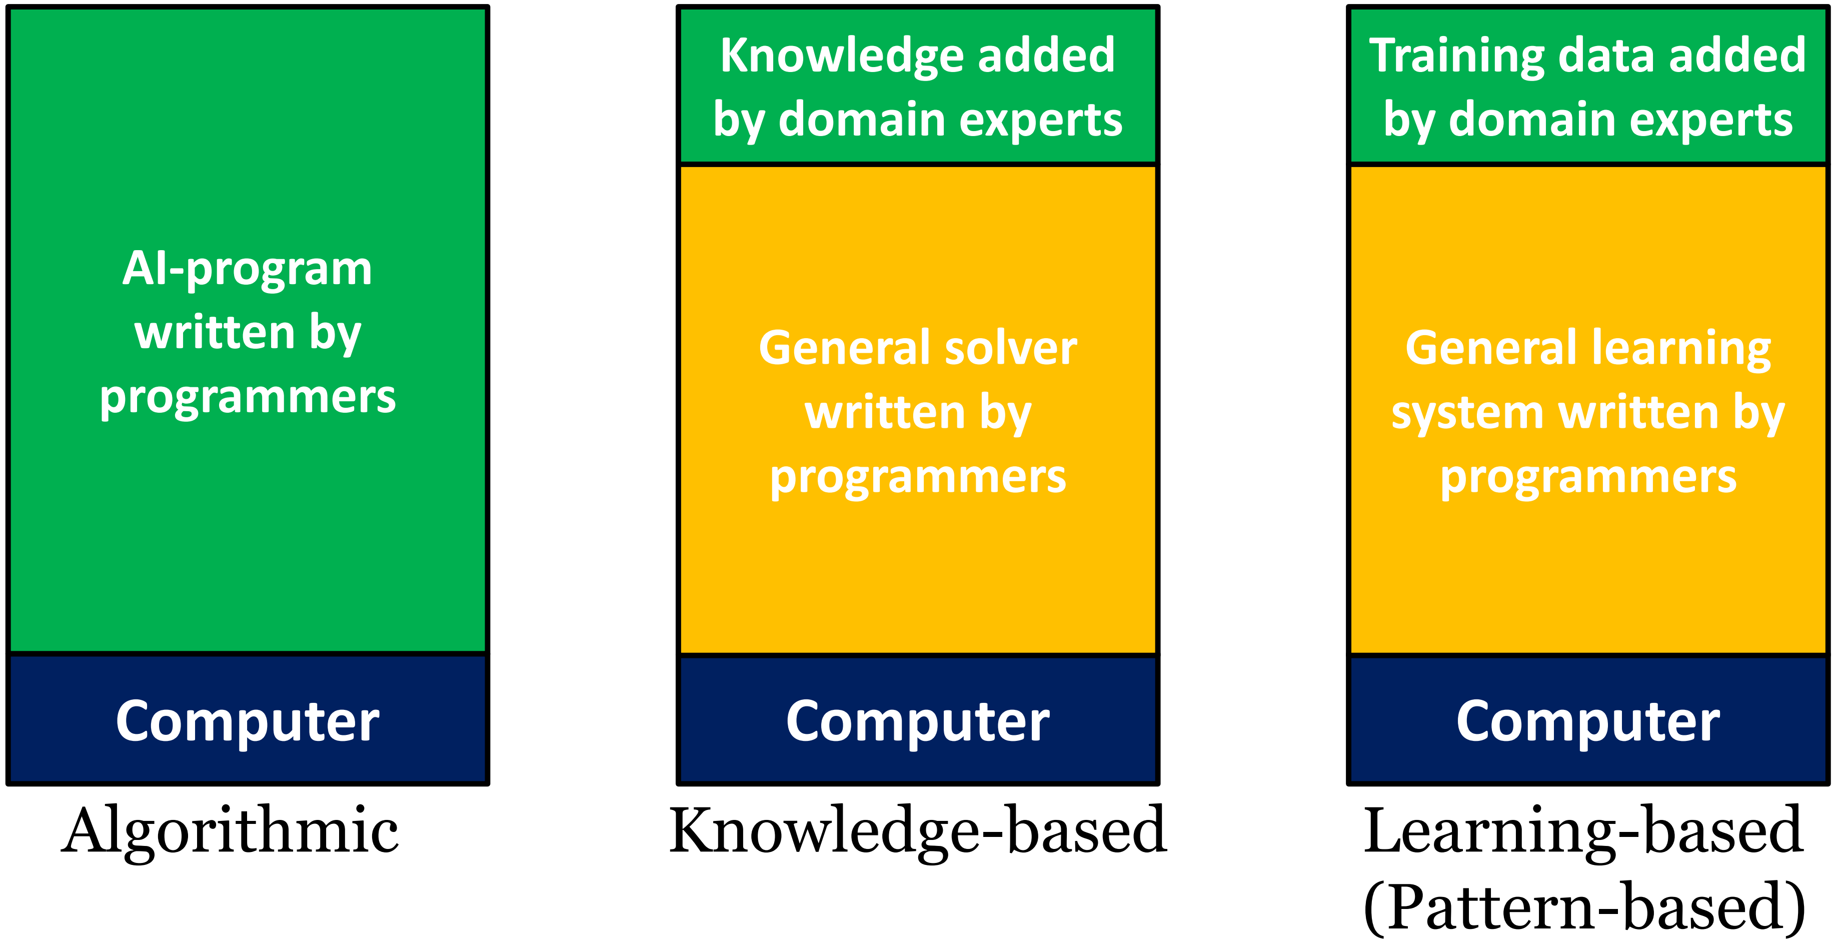

Representations and Search

Search

- Uninformed search
  - Depth-first search
  - Breadth-first search
  - Recursive backtracking
- Informed / Heuristic search
  - Best-first search
  - A* search
  - Branch-and-bound
  - Hill-climbing / Gradient descent
- Stochastic Search
  - Monte Carlo Tree Search
  - Stochastic Gradient Descent

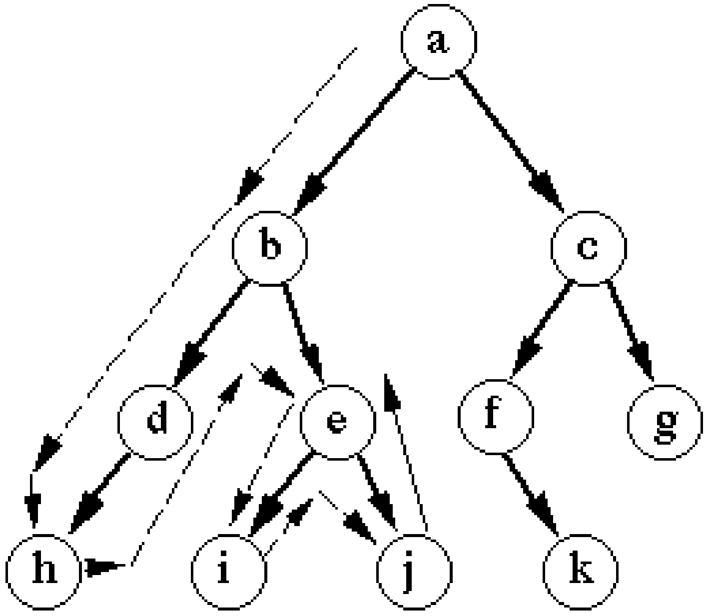

Depth-first search

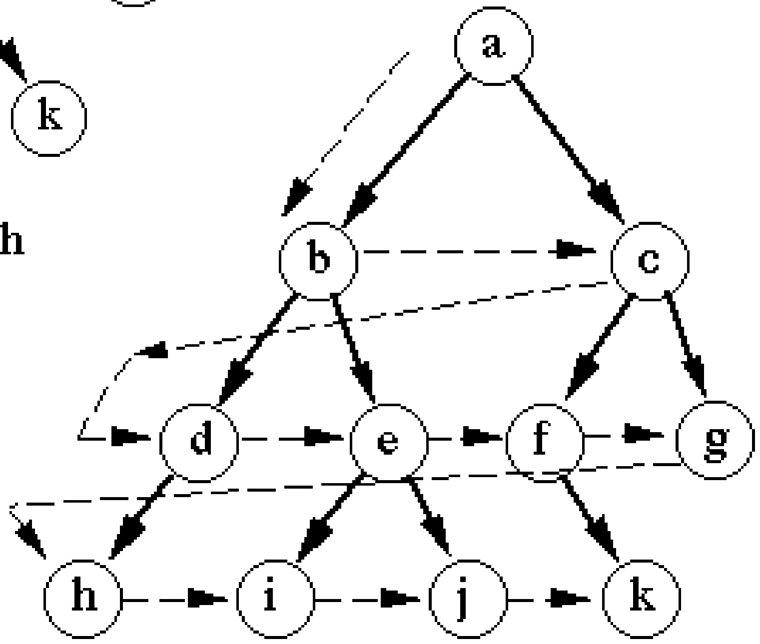

Breadth-first search

Search – Problem Definition

- **Initial State** : The state in which the agent starts or initial condition of the agent.
- **States** : All states that are reachable from initial state by any sequence of actions or all possible states that the agent can take. This is also referred to as State space.
- **Actions** : All possible actions that the agent can execute. Specifically, it provides the list of actions, that an agent can perform in a particular state. This is also referred to as Action space.
- **Transition Model** : This property describes the results of each action taken in a particular state.
- **Goal Test** : A way to check, whether a state is the goal.
- **Path Cost** : A function that assigns a numeric cost to a path w.r.t. performance measure

Search – Search Space

- **State space**: physical configuration
- **Search space**: abstract configuration often represented by a search tree or graph where a path is a possible solution.
- **Search tree**: representation of configurations and how they are connected by actions. A path represents a sequence of actions. The *root* is the initial state. The actions taken make the *branches* and the *nodes* are results of those actions. A node has depth, path cost and associated state in the state space.

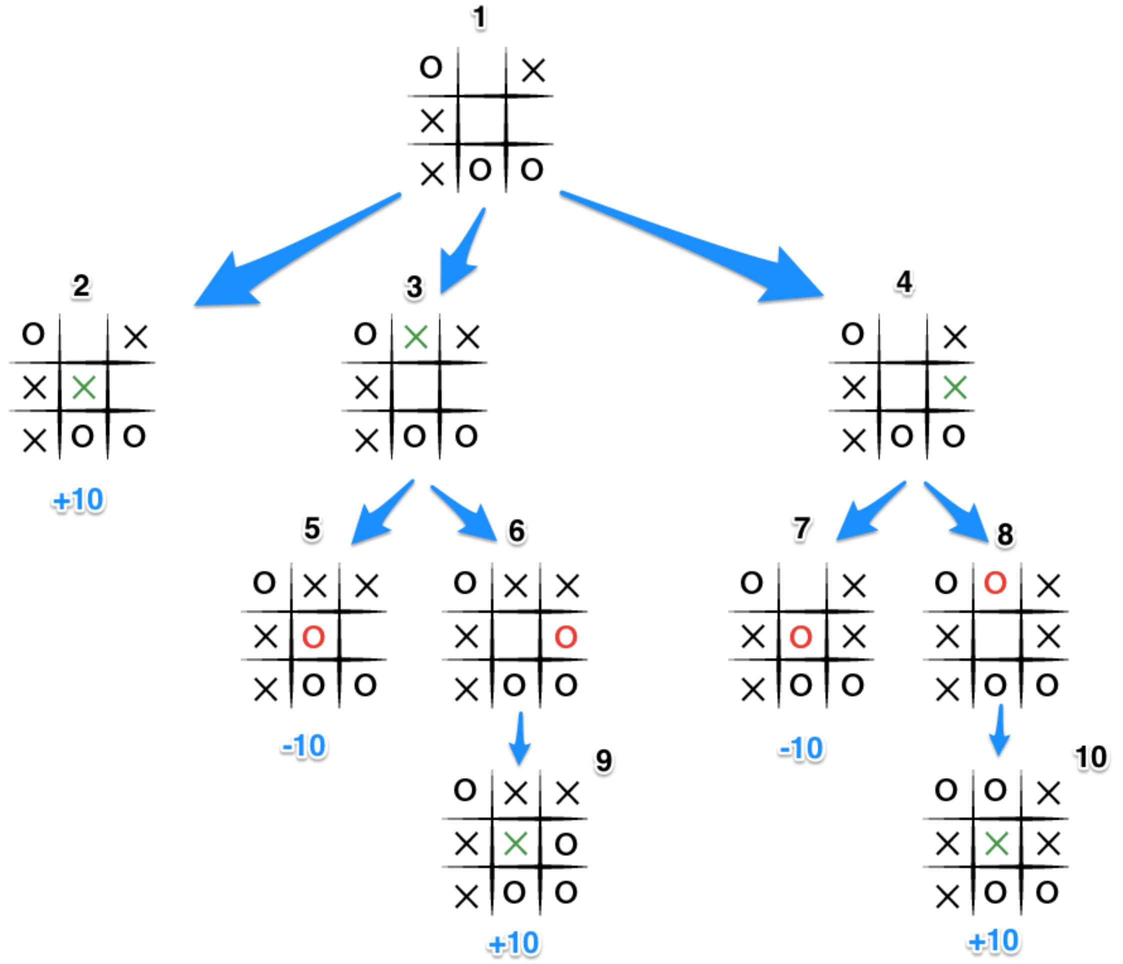

Mini-Max

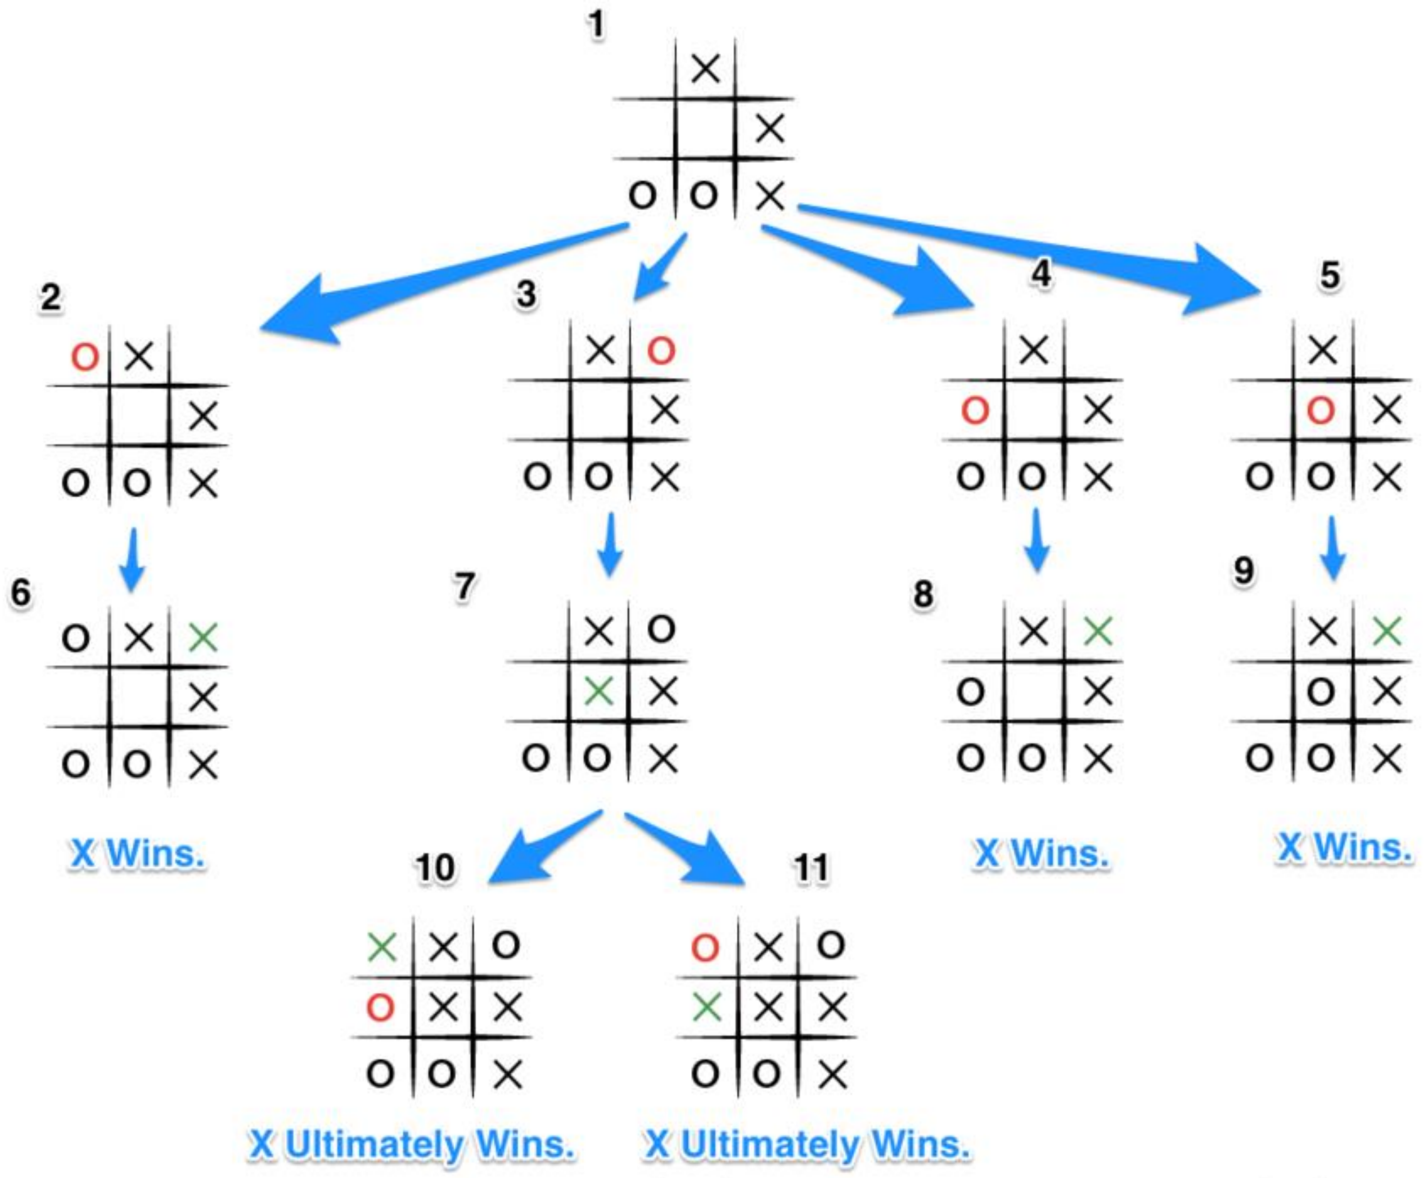

https://www.neverstopbuilding.com/blog/minimax

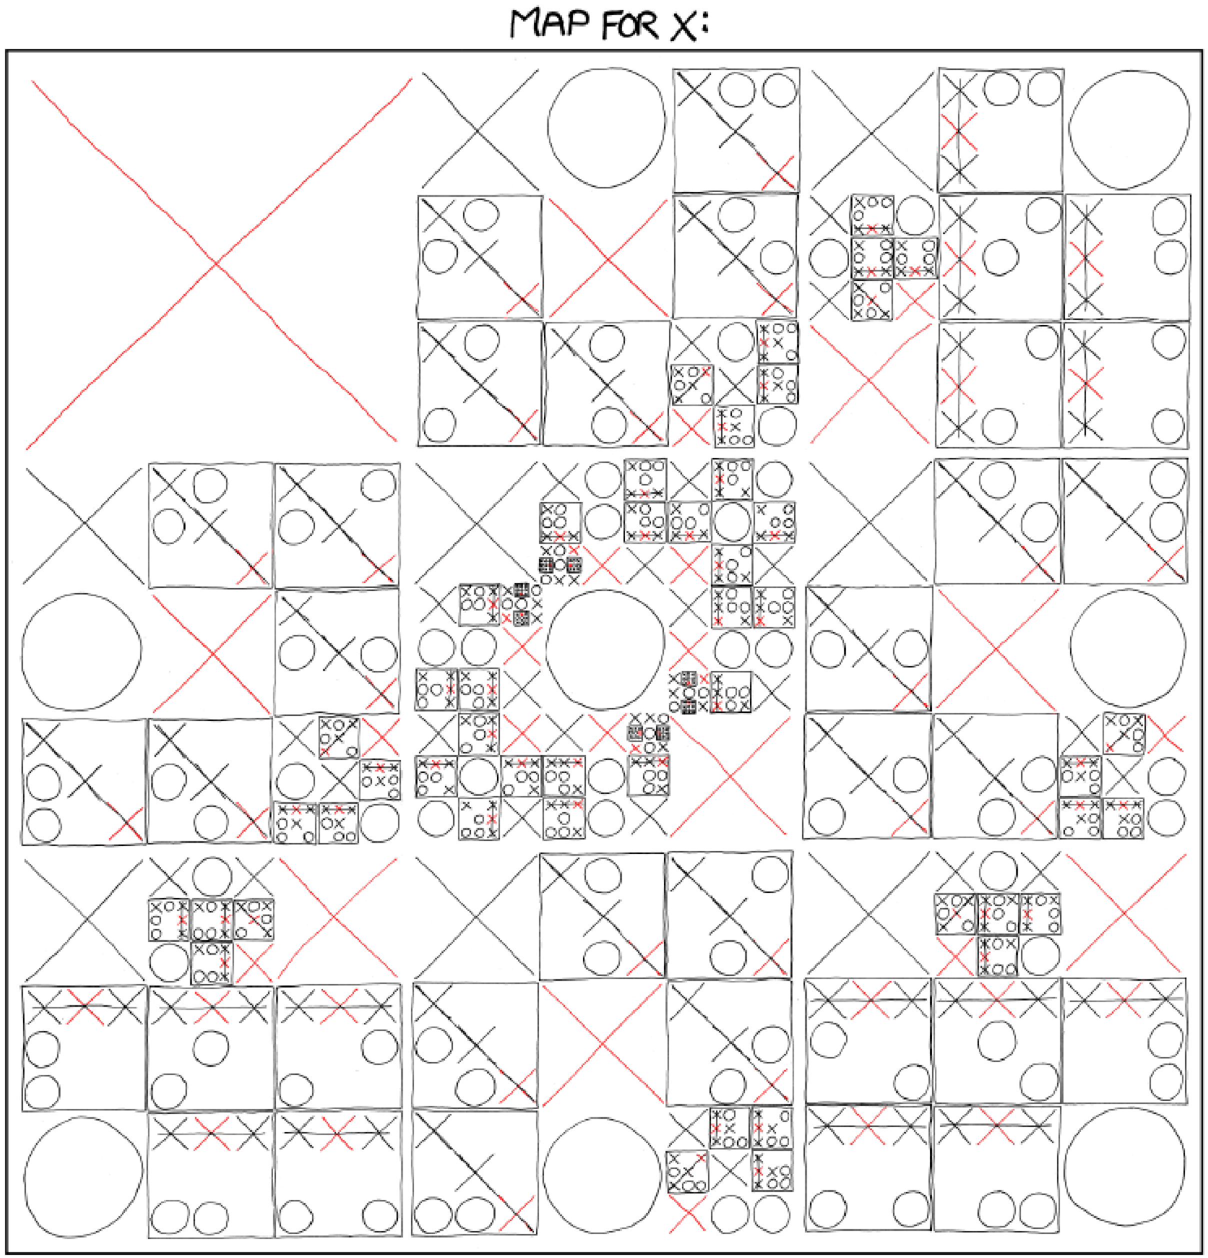

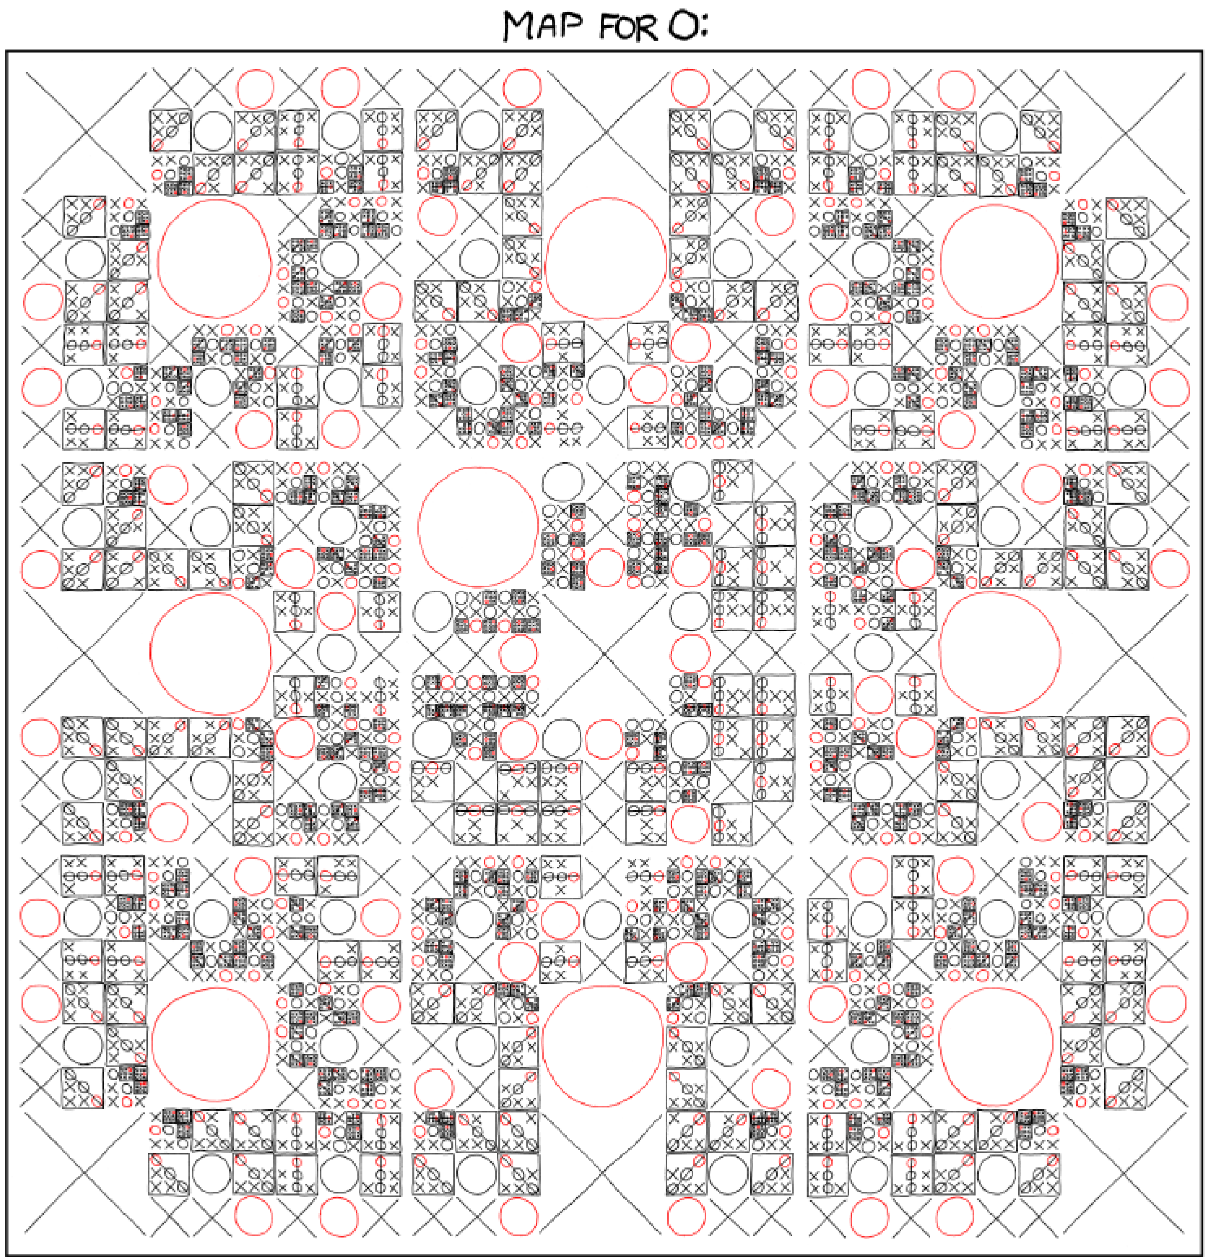

http://xkcd.com/832/

Efficiency and Representation

- Representations can be analyzed and their efficiency proved.
- The choice of abstraction / representation influences efficiency

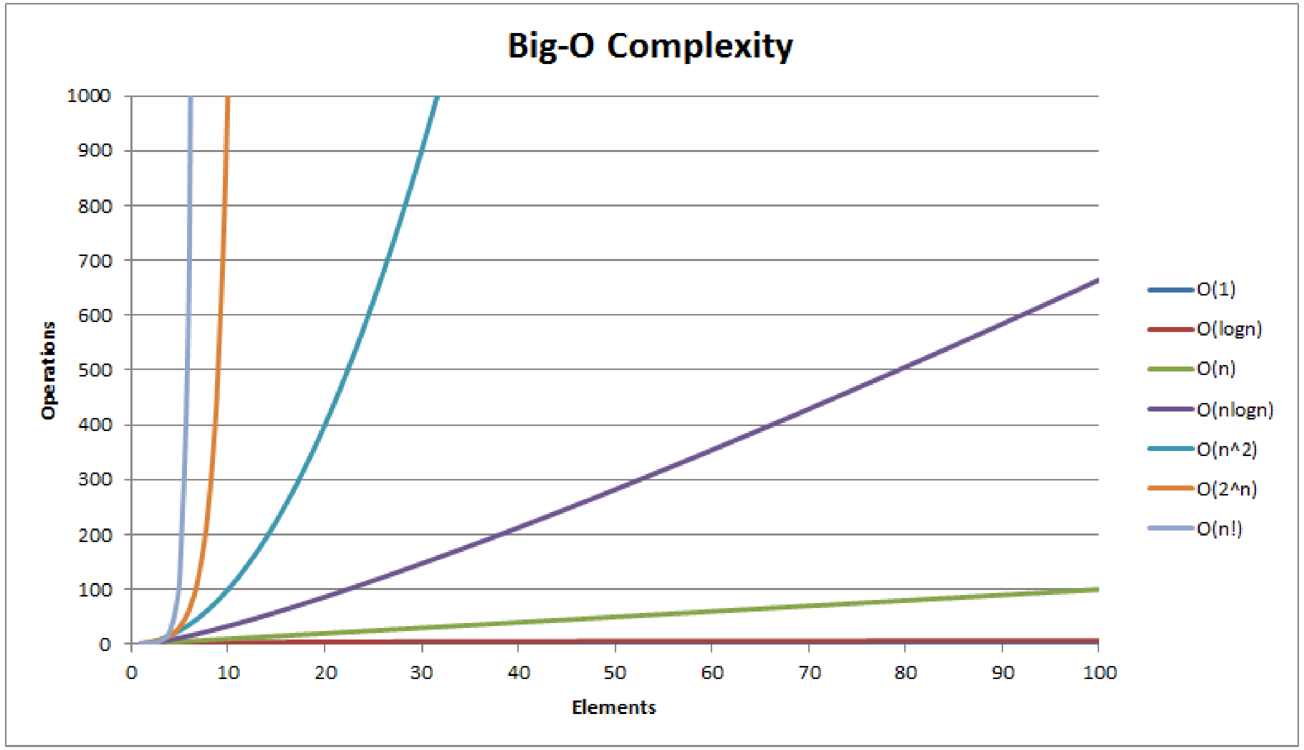

Graph Operations

Node / Edge Management	Storage	Add Vertex	Add Edge	Remove Vertex	Remove Edge	Query
Adjacency list	O(|V|+|E|)	O(1)	O(1)	O(|V| + |E|)	O(|E|)	O(|V|)
Incidence list	O(|V|+|E|)	O(1)	O(1)	O(|E|)	O(|E|)	O(|E|)
Adjacency matrix	O(|V|^2)	O(|V|^2)	O(1)	O(|V|^2)	O(1)	O(1)
Incidence matrix	O(|V| · |E|)	O(|V| · |E|)	O(|V| · |E|)	O(|V| · |E|)	O(|V| · |E|)	O(|E|)

Excellent Good Fair

Bad Horrible

Applications of Search

- Game playing (chess, Go, …)
- Constraint satisfaction
- Optimization
- Machine learning
- Planning
- …

Knowledge Representation and Reasoning (KR&R)

- Knowledge representation and reasoning is a major sub-area of AI.
- *Intelligence can be understood by studying knowledge.*
- *Knowledge* is often defined as true justified belief in epistemology.
  - Declarative knowledge
  - Procedural knowledge
  - Heuristic knowledge
- *Representation* is a relationship between two domains, where the first is meant to "stand for" or take the place of the second.
- *Reasoning* is the formal manipulation of symbols representing knowledge to produce a new set of symbols representing new knowledge.

The Physical-Symbol System Hypothesis

- A physical-symbol system has the necessary and sufficient means for general intelligent action.
  - *Necessary*: any system exhibiting intelligence will prove upon analysis to be a physical symbol system.
  - *Sufficient:* any physical-symbol system of sufficient size can be organized further to exhibit general intelligence.

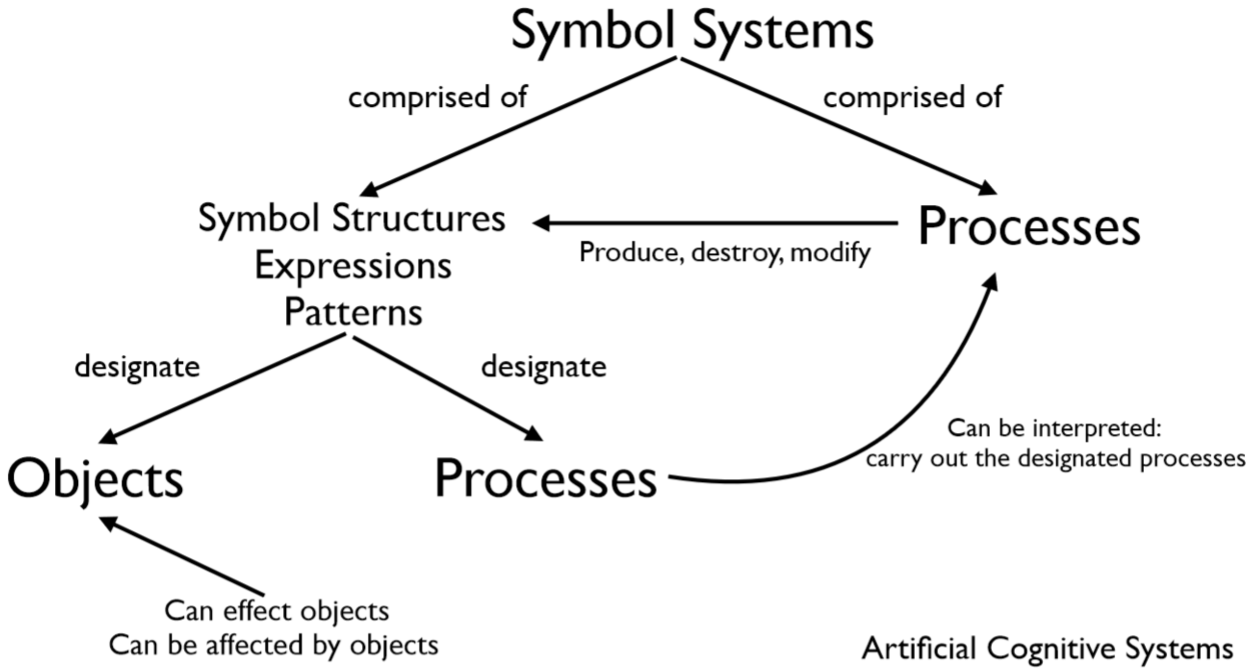

Artificial Cognitive Systems
David Vernon, MIT Press

Newell, Allen; Simon, H. A. (1976), "Computer Science as Empirical Inquiry: Symbols and Search", Communications of the ACM, 19 (3): 113–126

Machine Learning

Machine Learning

- Machine Learning is a branch of artificial intelligence that provides the computer system the ability to progressively learn and improve its performance on handling various tasks without being explicitly programmed to perform all the task.
- Another definition of Machine Learning explains it as: the process of trying to deduce unknown values from known values.
- More formally, study of algorithms that
  - improve their performance P
  - at some task T
  - with experience E

Machine Learning

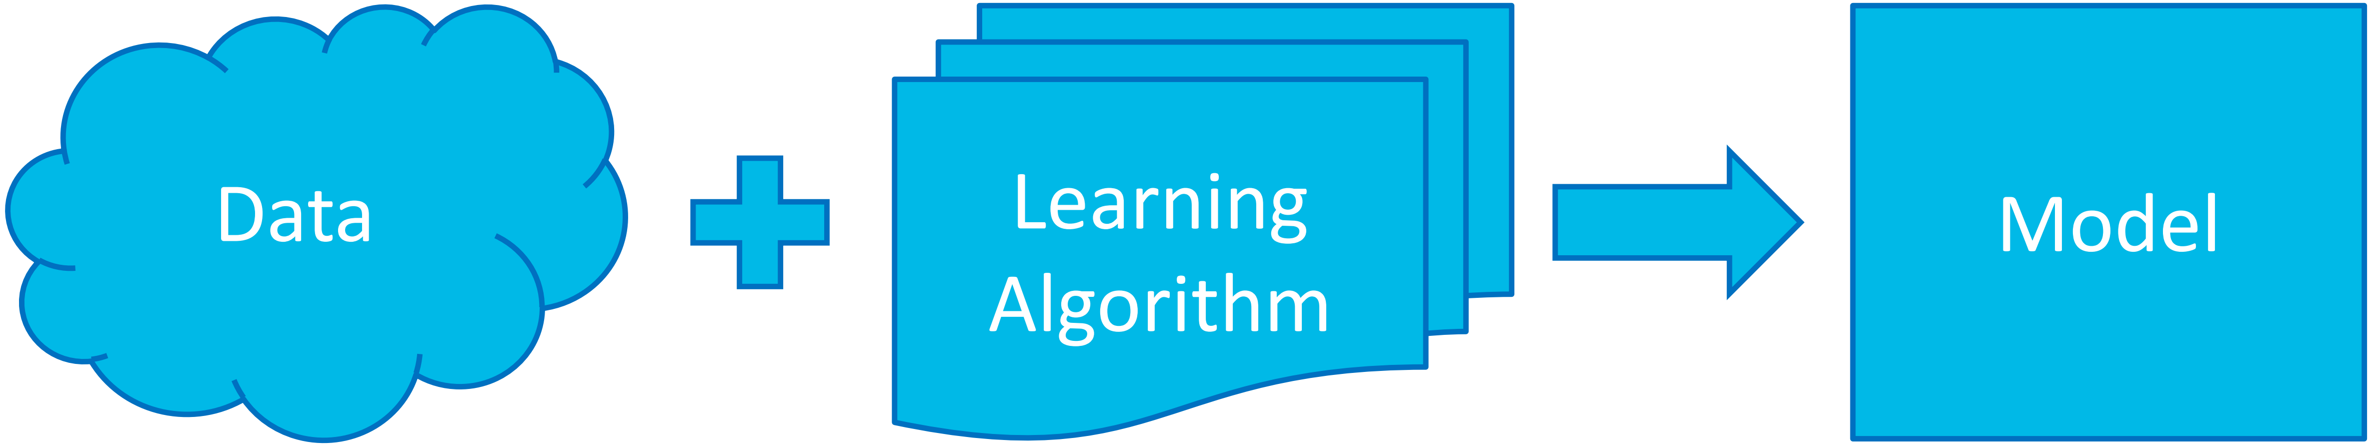

Machine Learning

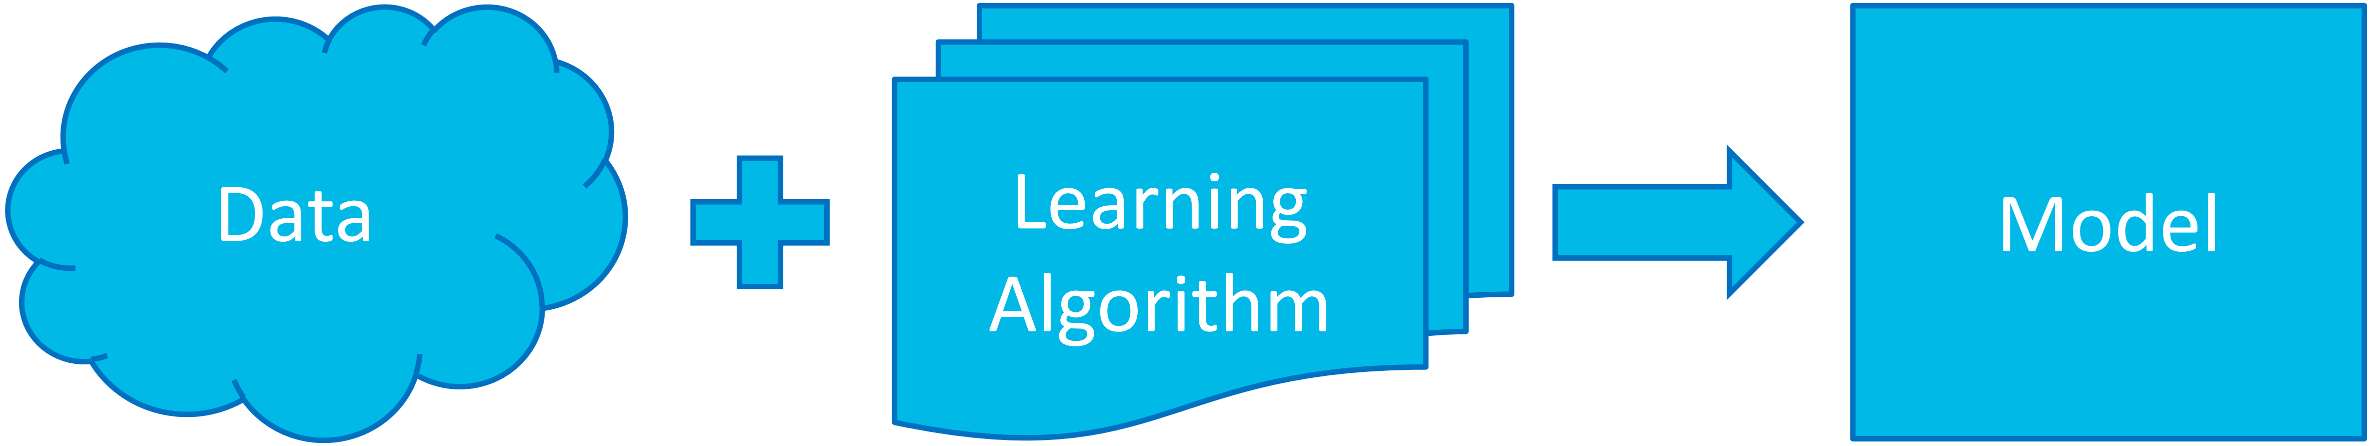

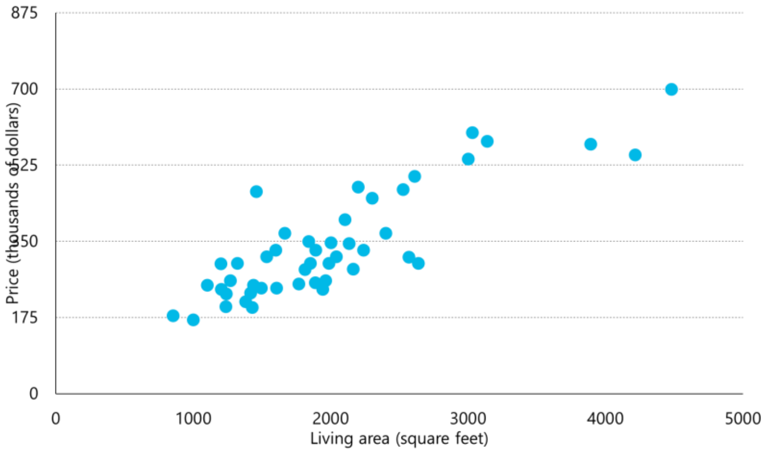

Linear regression

$$\hat{y} = \boldsymbol{x\theta}$$

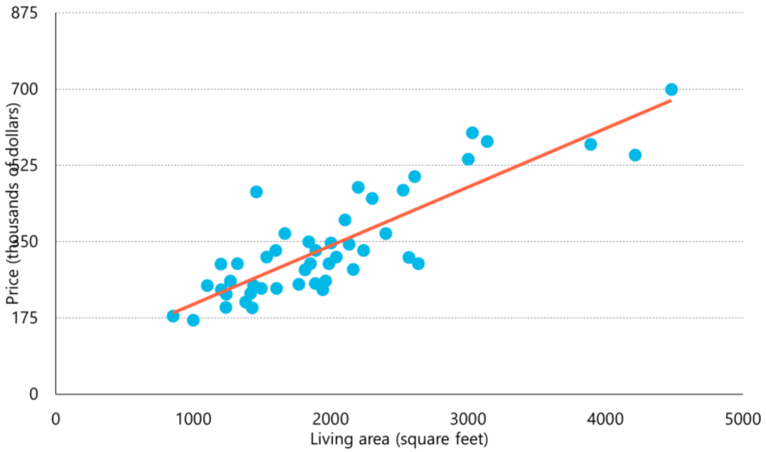

The Importance of Feature Selection

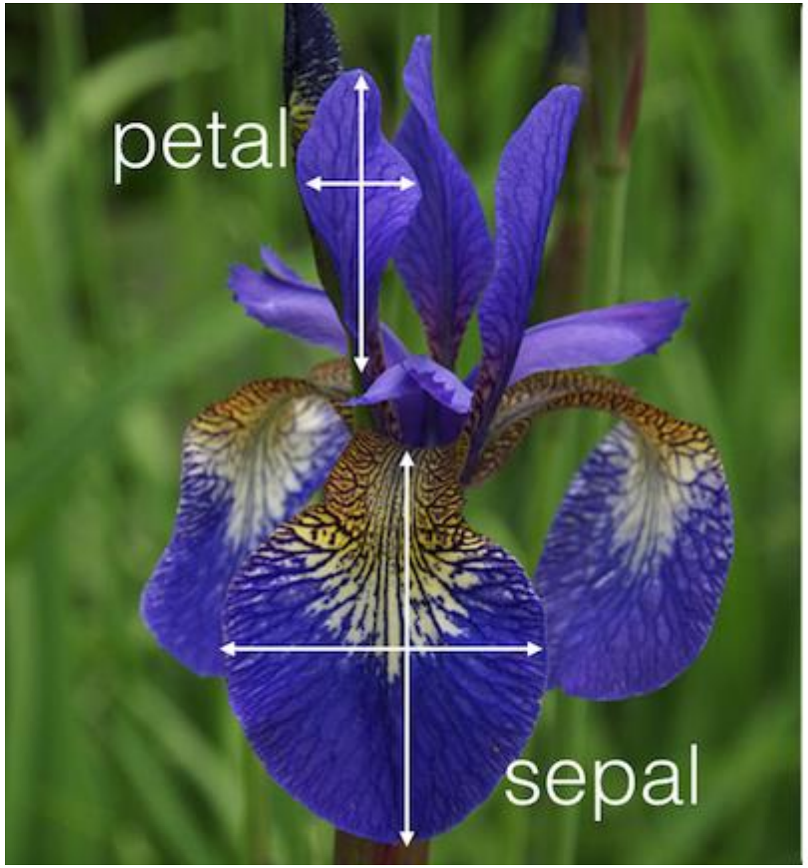

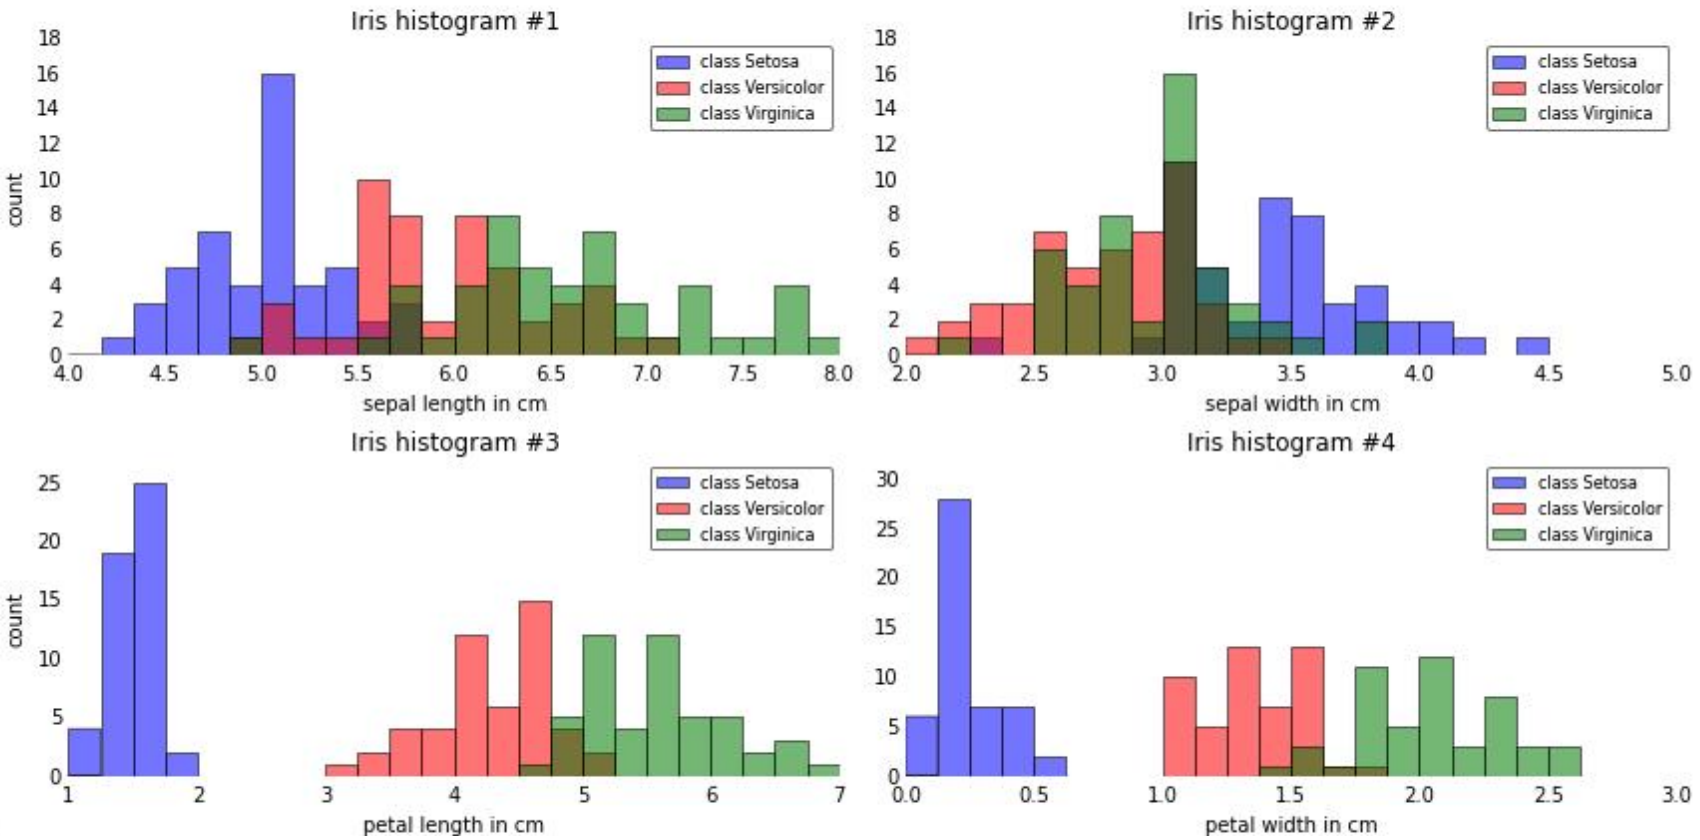

https://sebastianraschka.com/Articles/2014_intro_supervised_learning.html

Classification

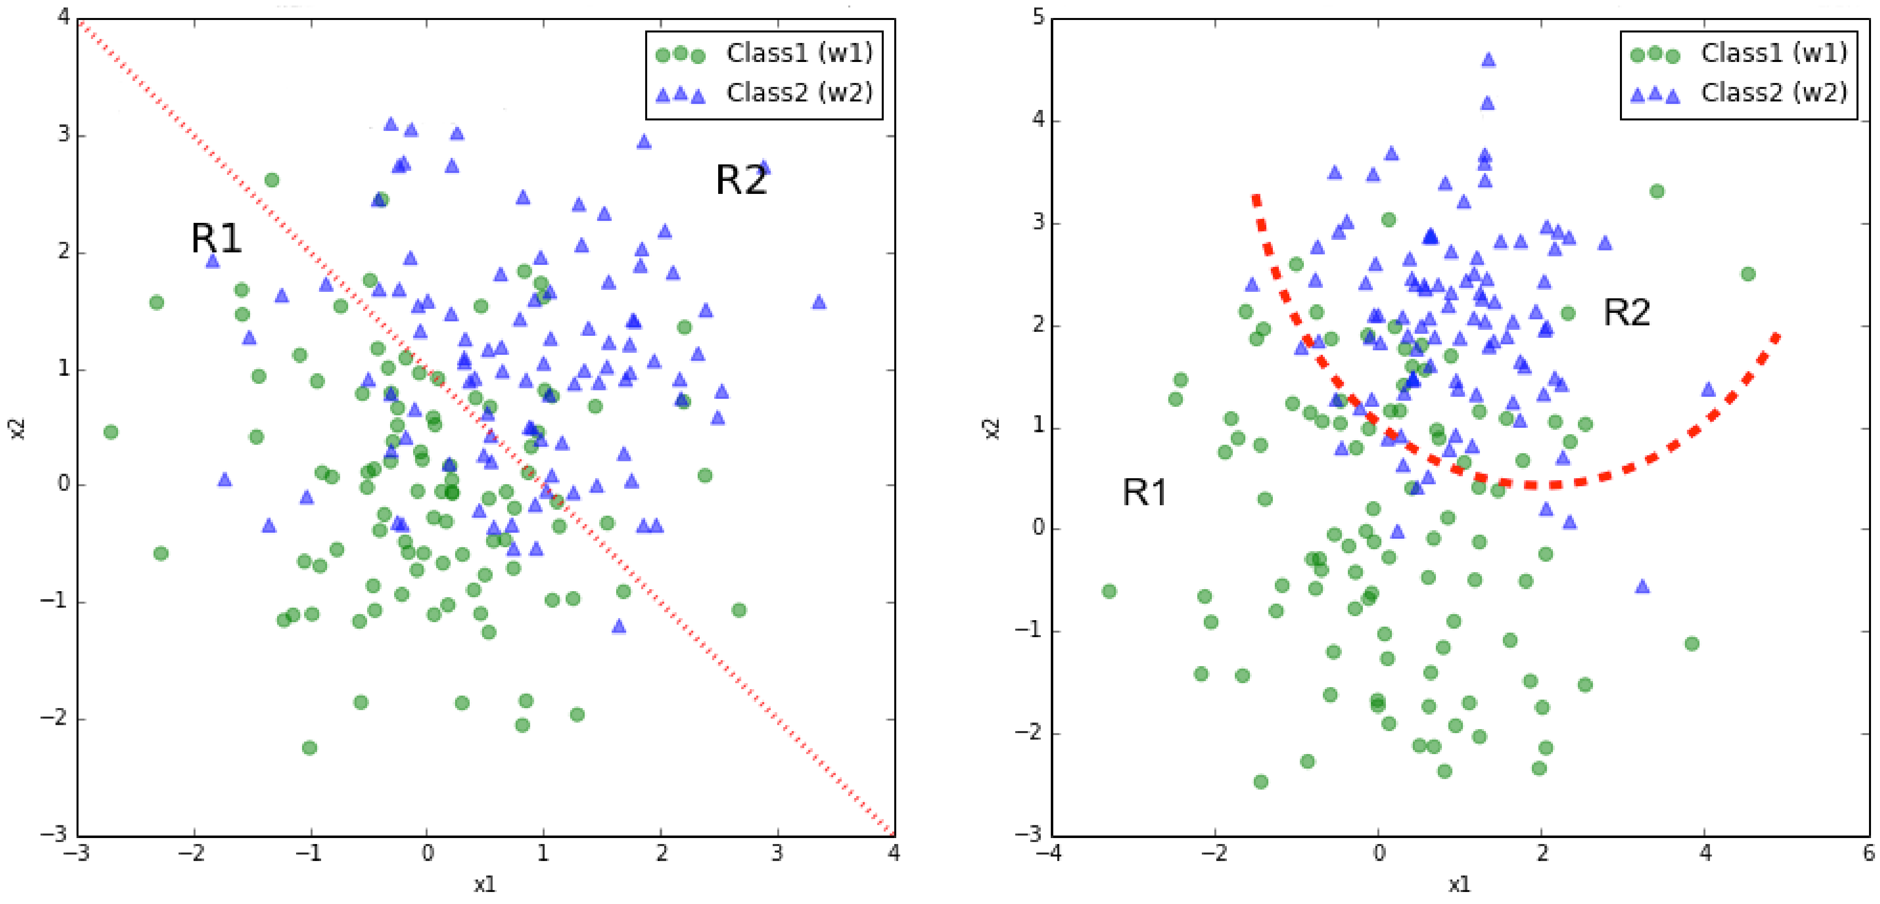

https://sebastianraschka.com/Articles/2014_intro_supervised_learning.html

Model Types

Classification

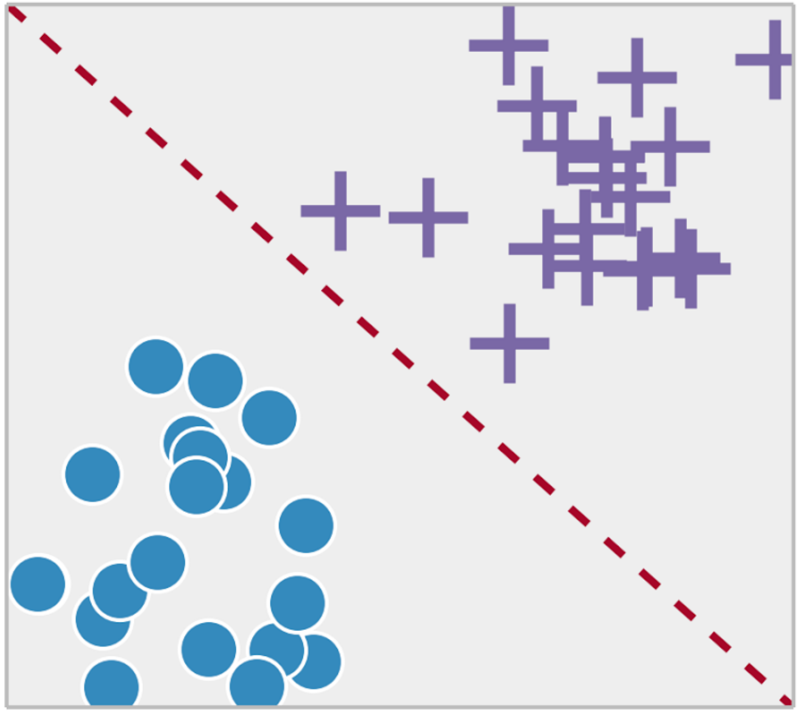

Regression

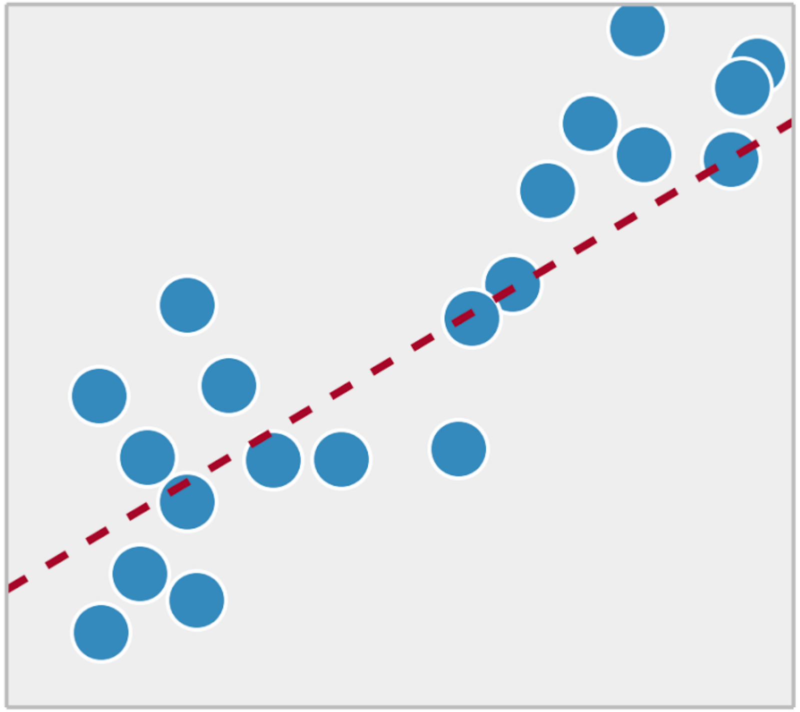

Generative

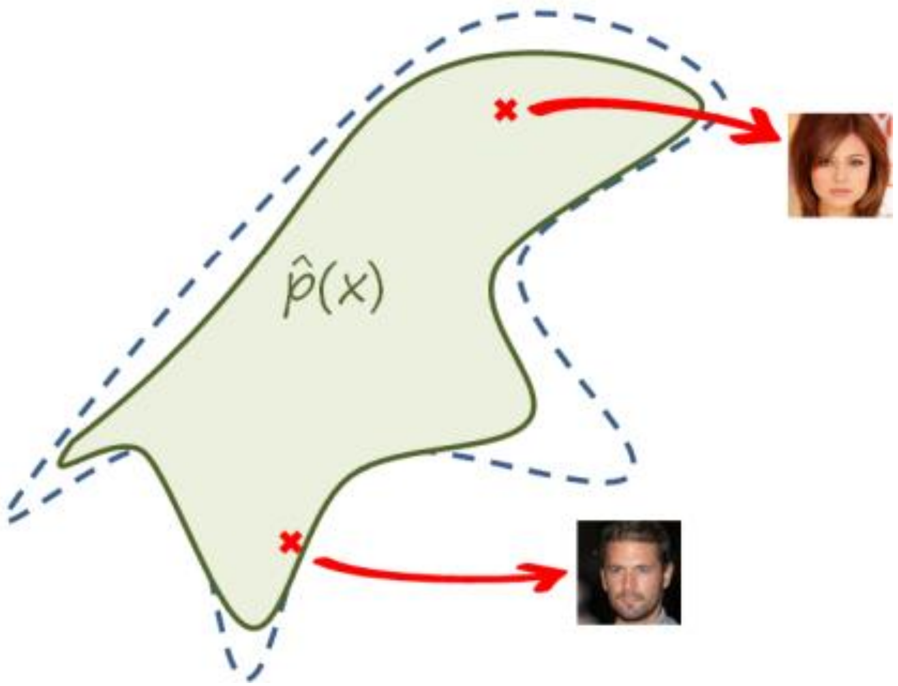

Types of Machine Learning

- Supervised learning
  - Given input-output examples f(X)=Y, learn the function f().
- Unsupervised learning
  - Given input examples, find patterns such as clusters
- Reinforcement learning
  - Select and execute an action, get feedback, update policy (what action to do in which state).

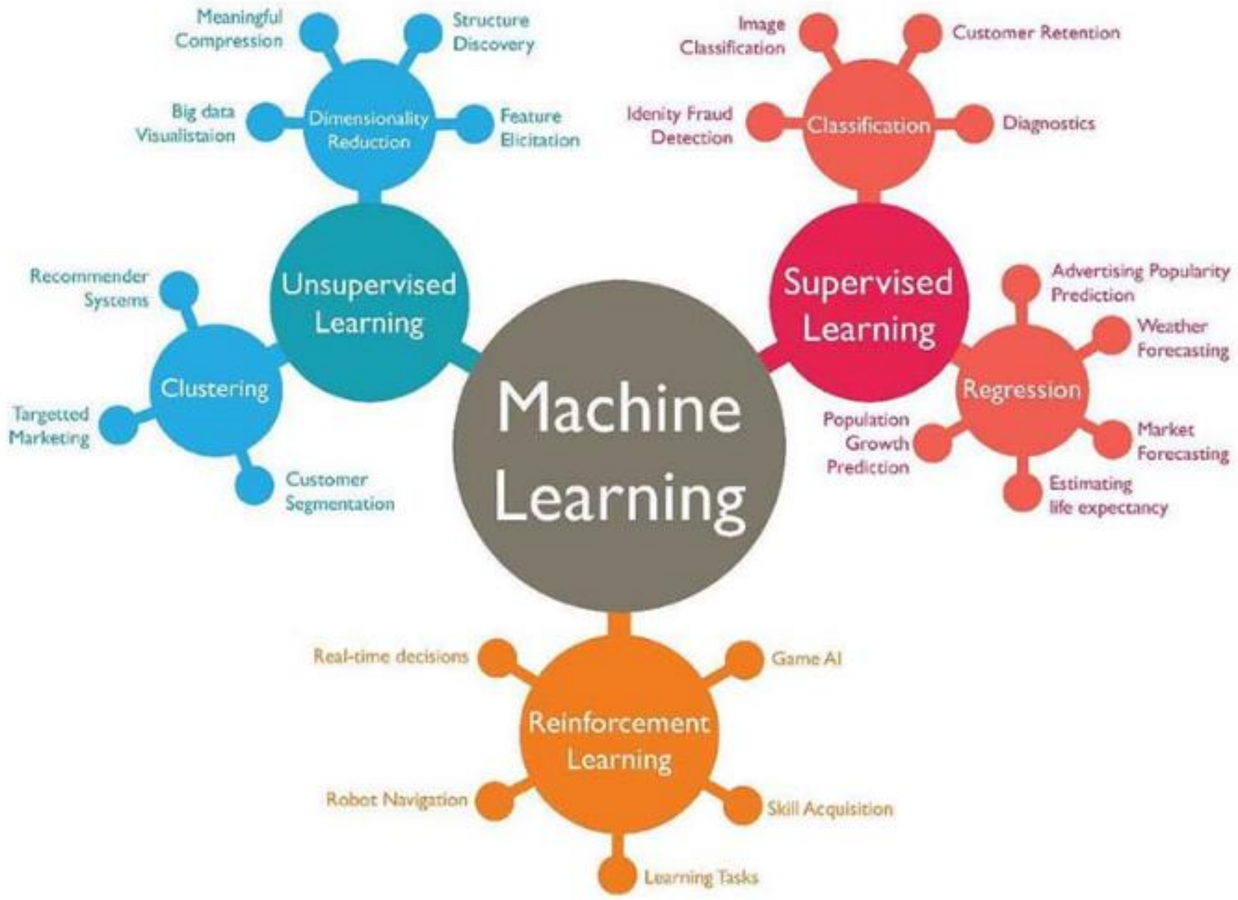

https://www.techleer.com/articles/203-machine-learning-algorithm-backbone-of-emerging-technologies/

Neural Networks

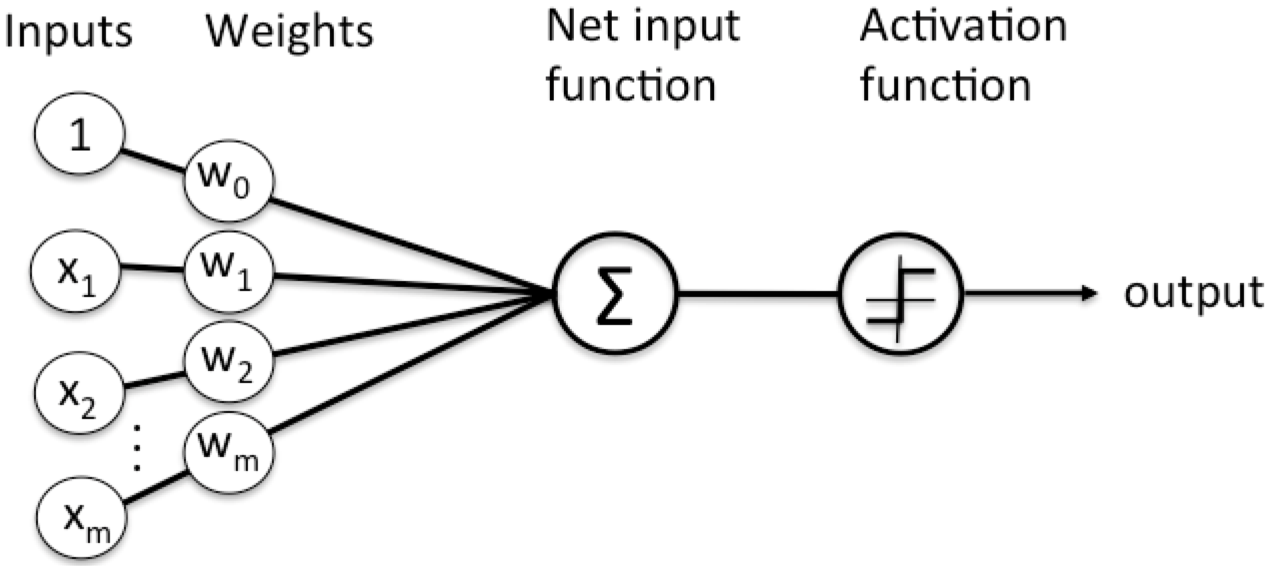

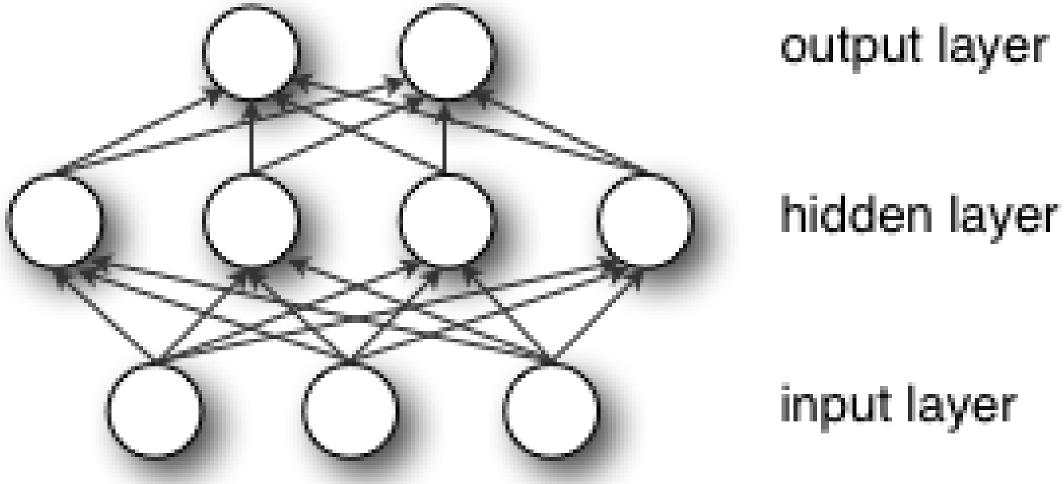

https://deeplearning4j.org/

Convolutional Neural Networks

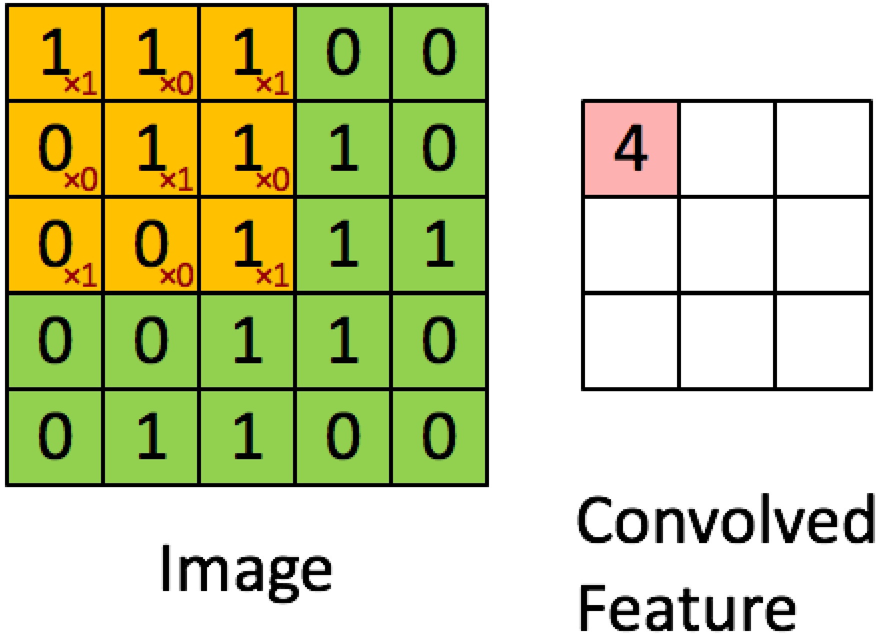

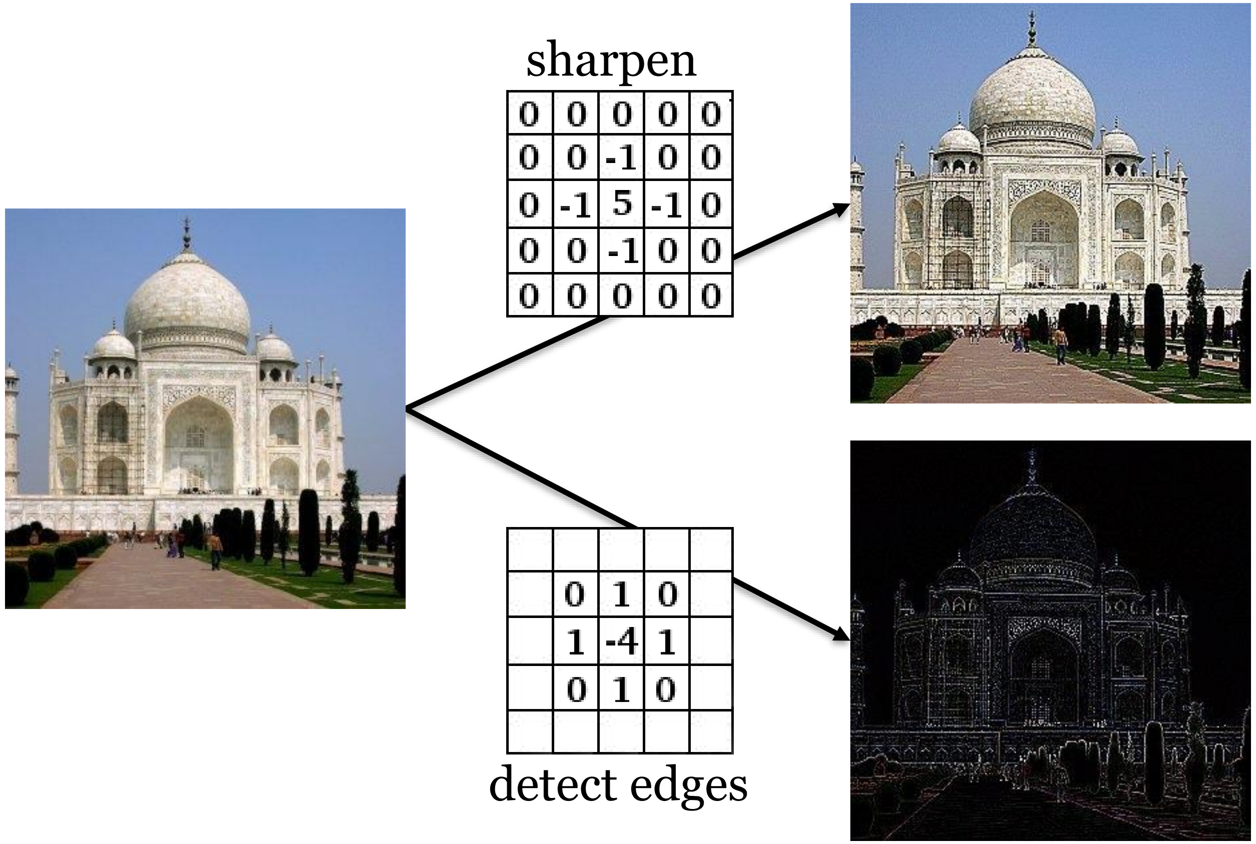

http://www.wildml.com/2015/11/understanding-convolutional-neural-networks-for-nlp/

Convolutional Neural Networks

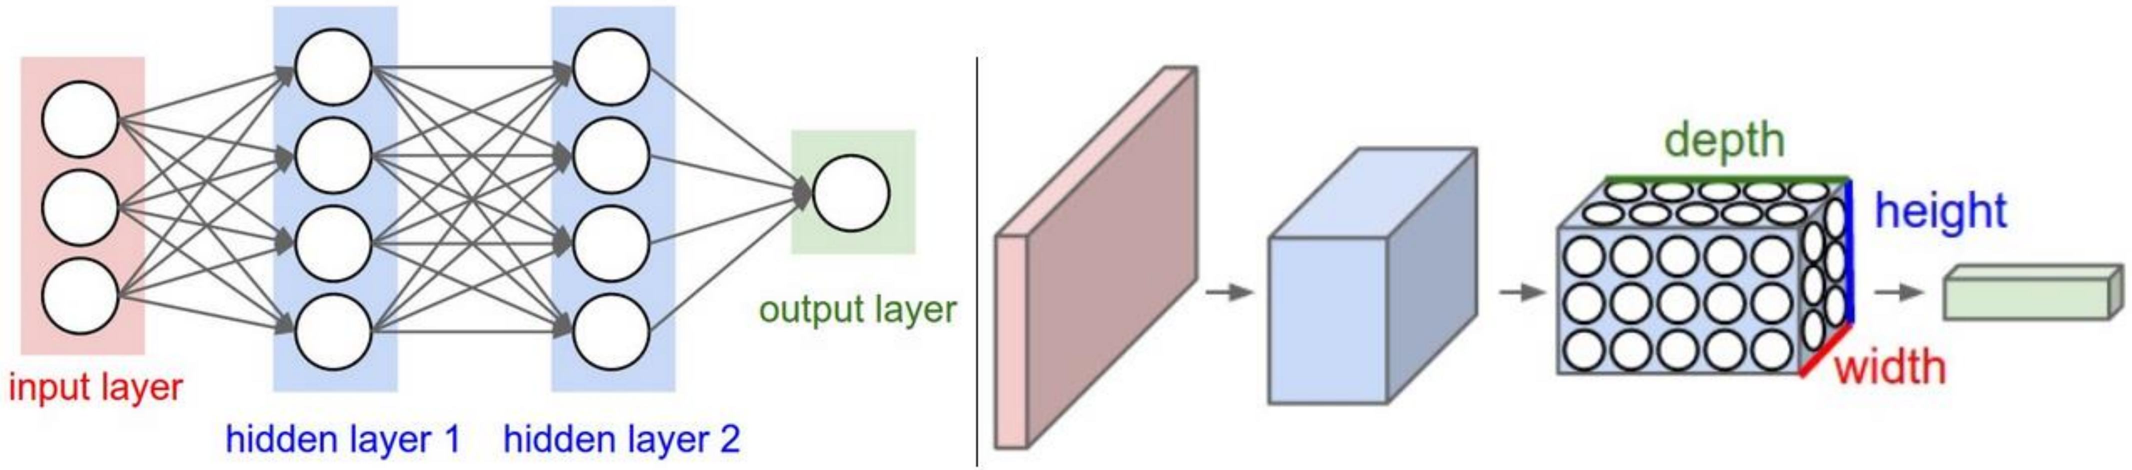

Left: A regular 3-layer Neural Network. Right: A ConvNet arranges its neurons in three dimensions (width, height, depth), as visualized in one of the layers. Every layer of a ConvNet transforms the 3D input volume to a 3D output volume of neuron activations. In this example, the red input layer holds the image, so its width and height would be the dimensions of the image, and the depth would be 3 (Red, Green, Blue channels).

http://cs231n.github.io/convolutional-networks/

Convolutional Neural Networks

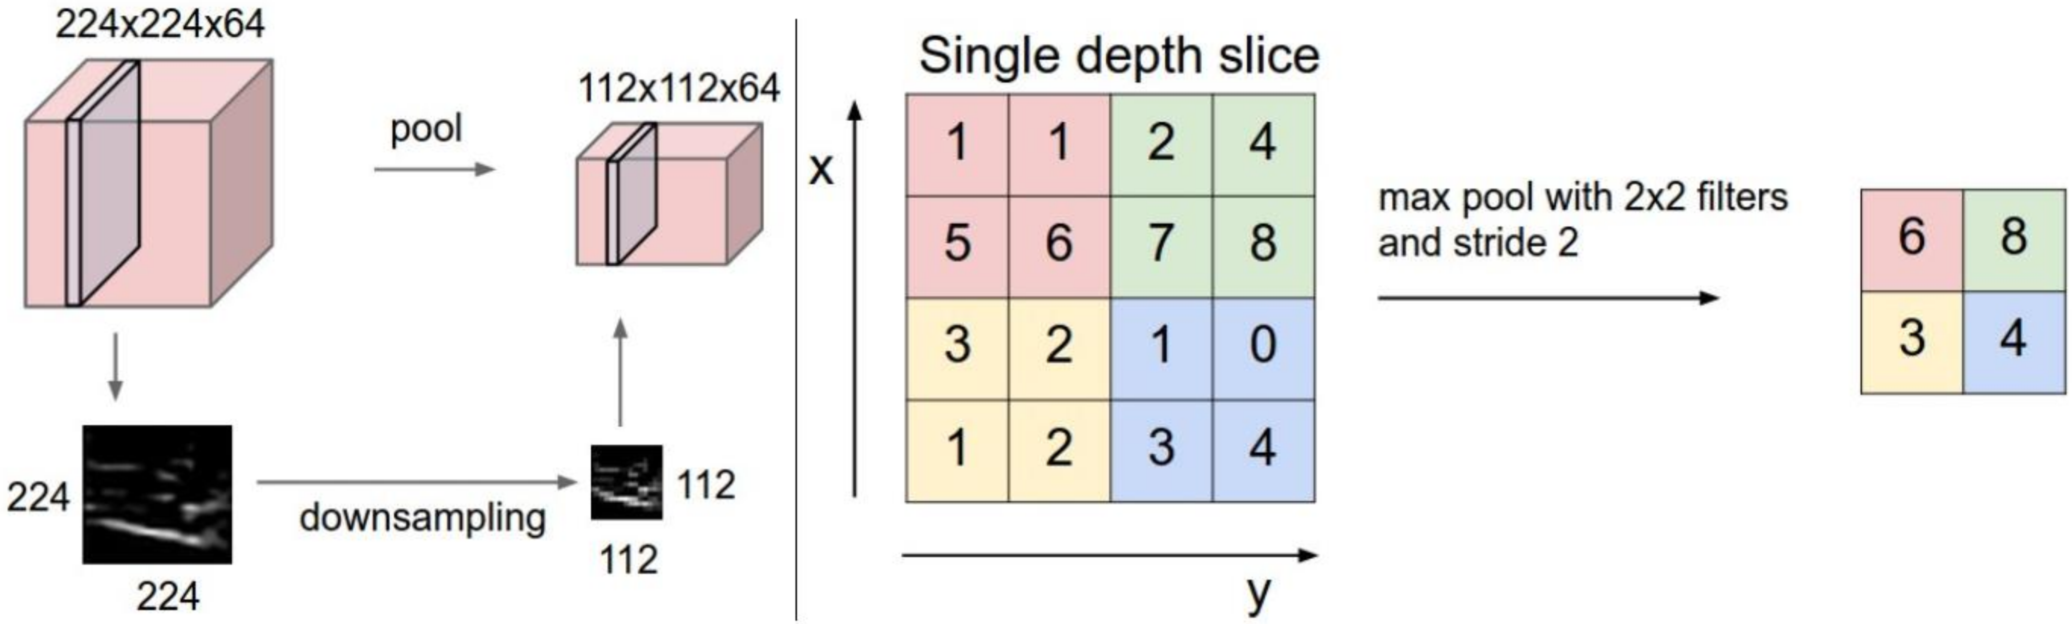

Pooling layer downsamples the volume spatially, independently in each depth slice of the input volume. **Left:** In this example, the input volume of size [224x224x64] is pooled with filter size 2, stride 2 into output volume of size [112x112x64]. Notice that the volume depth is preserved. **Right:** The most common downsampling operation is max, giving rise to **max pooling**, here shown with a stride of 2. That is, each max is taken over 4 numbers (little 2x2 square).

http://cs231n.github.io/convolutional-networks/

Deep Neural Networks

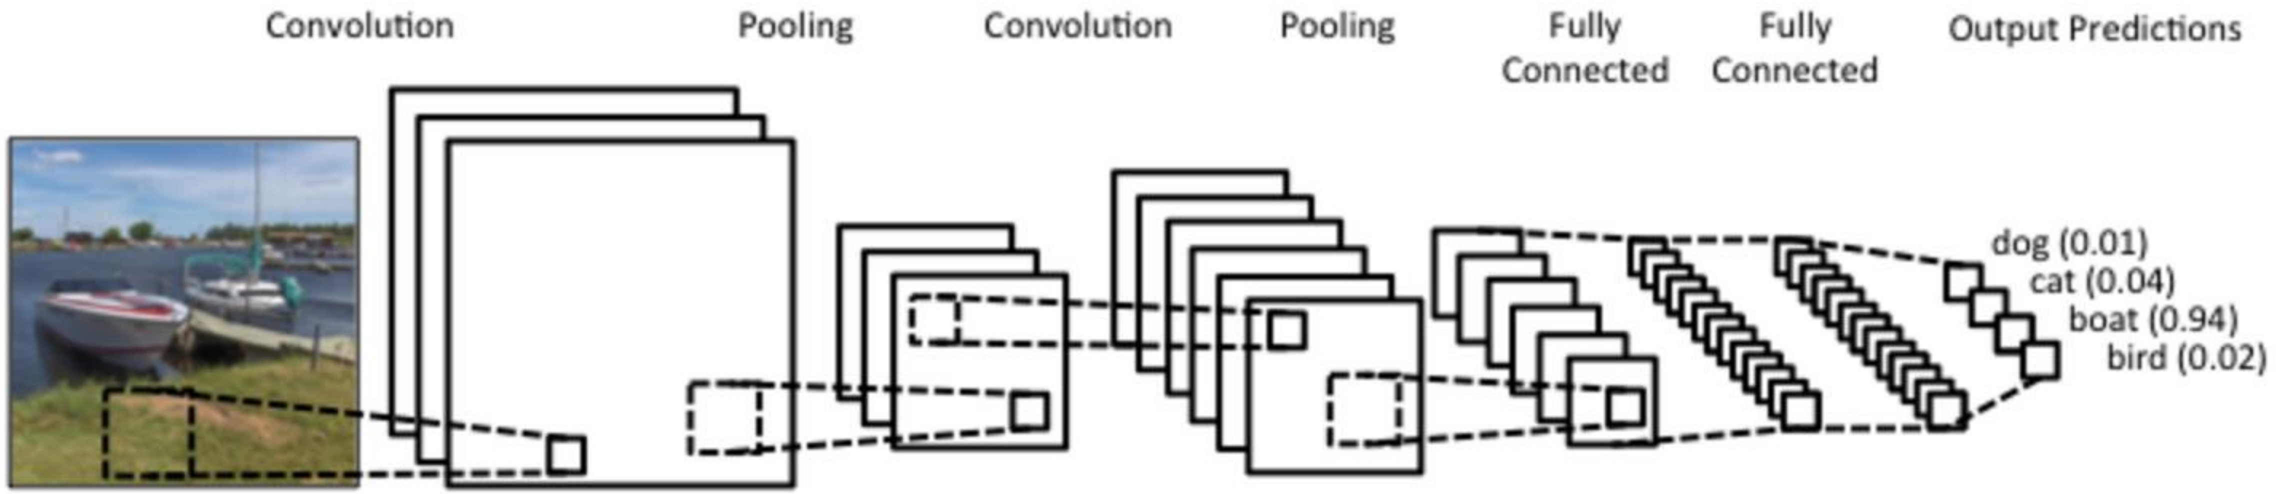

http://www.wildml.com/2015/11/understanding-convolutional-neural-networks-for-nlp/

Deep Neural Networks

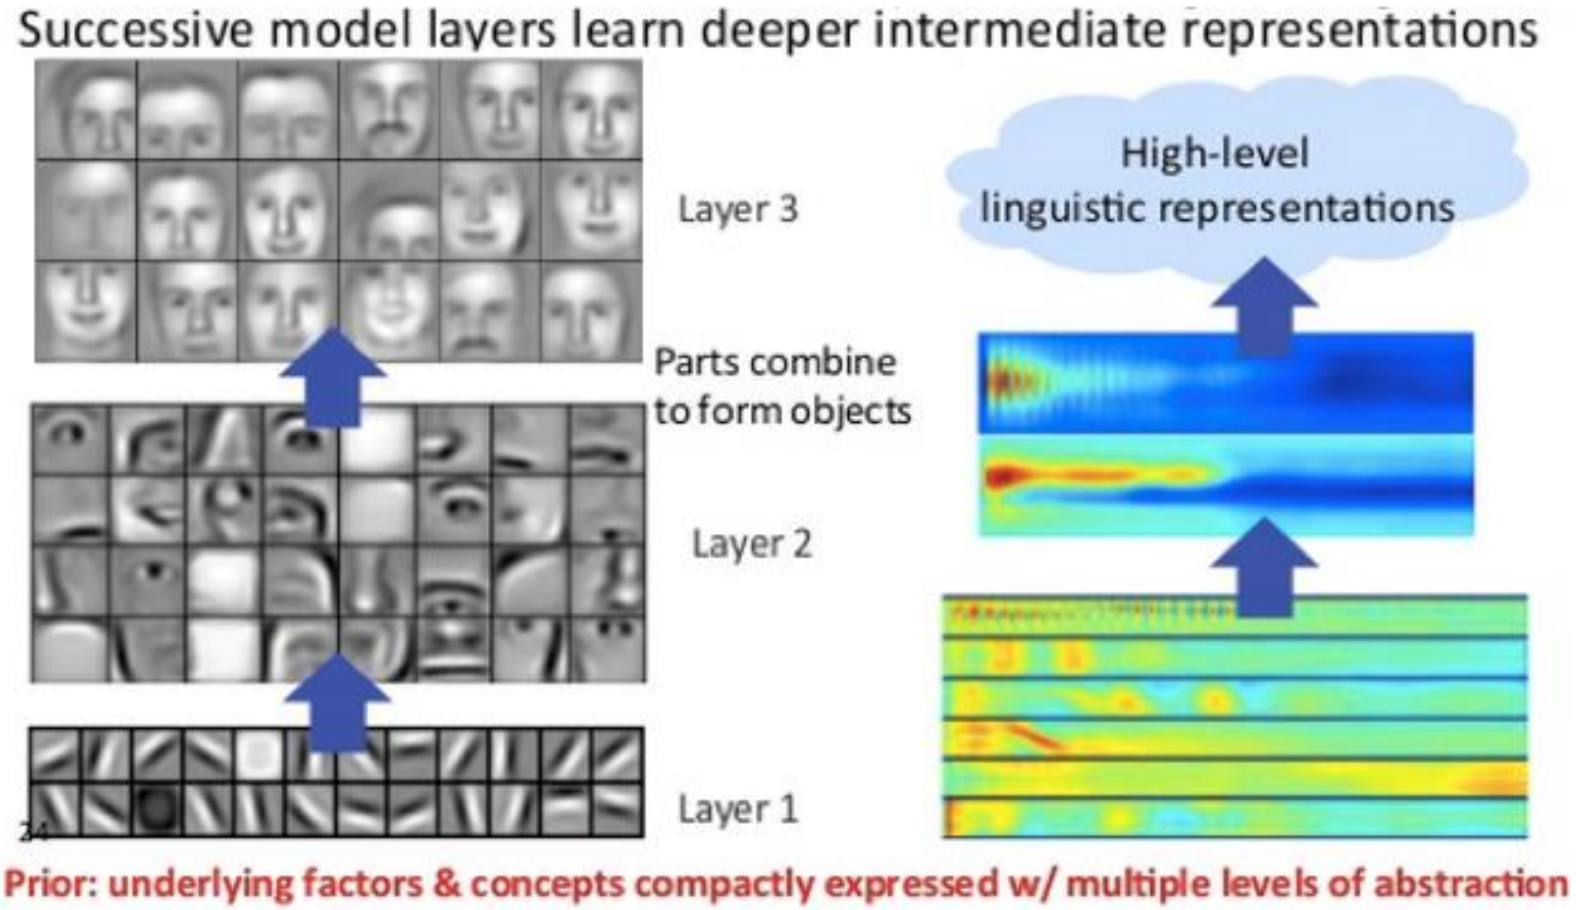

Recurrent Neural Networks

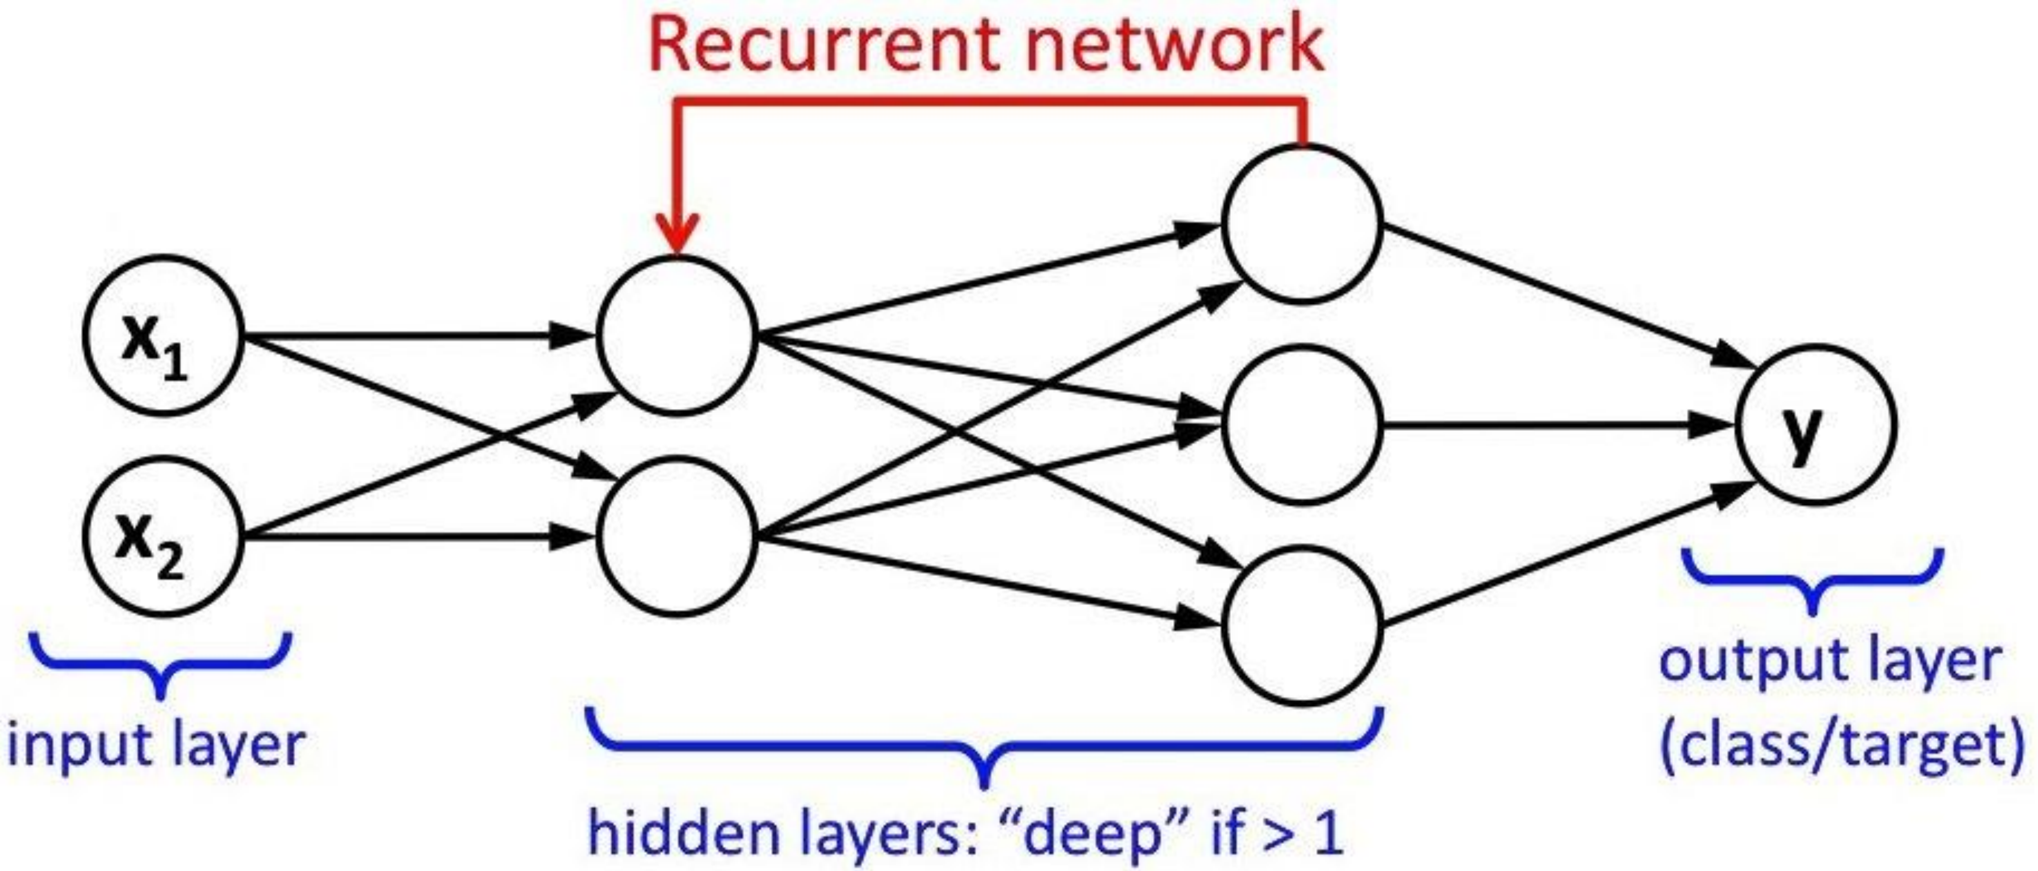

Recurrent Neural Networks

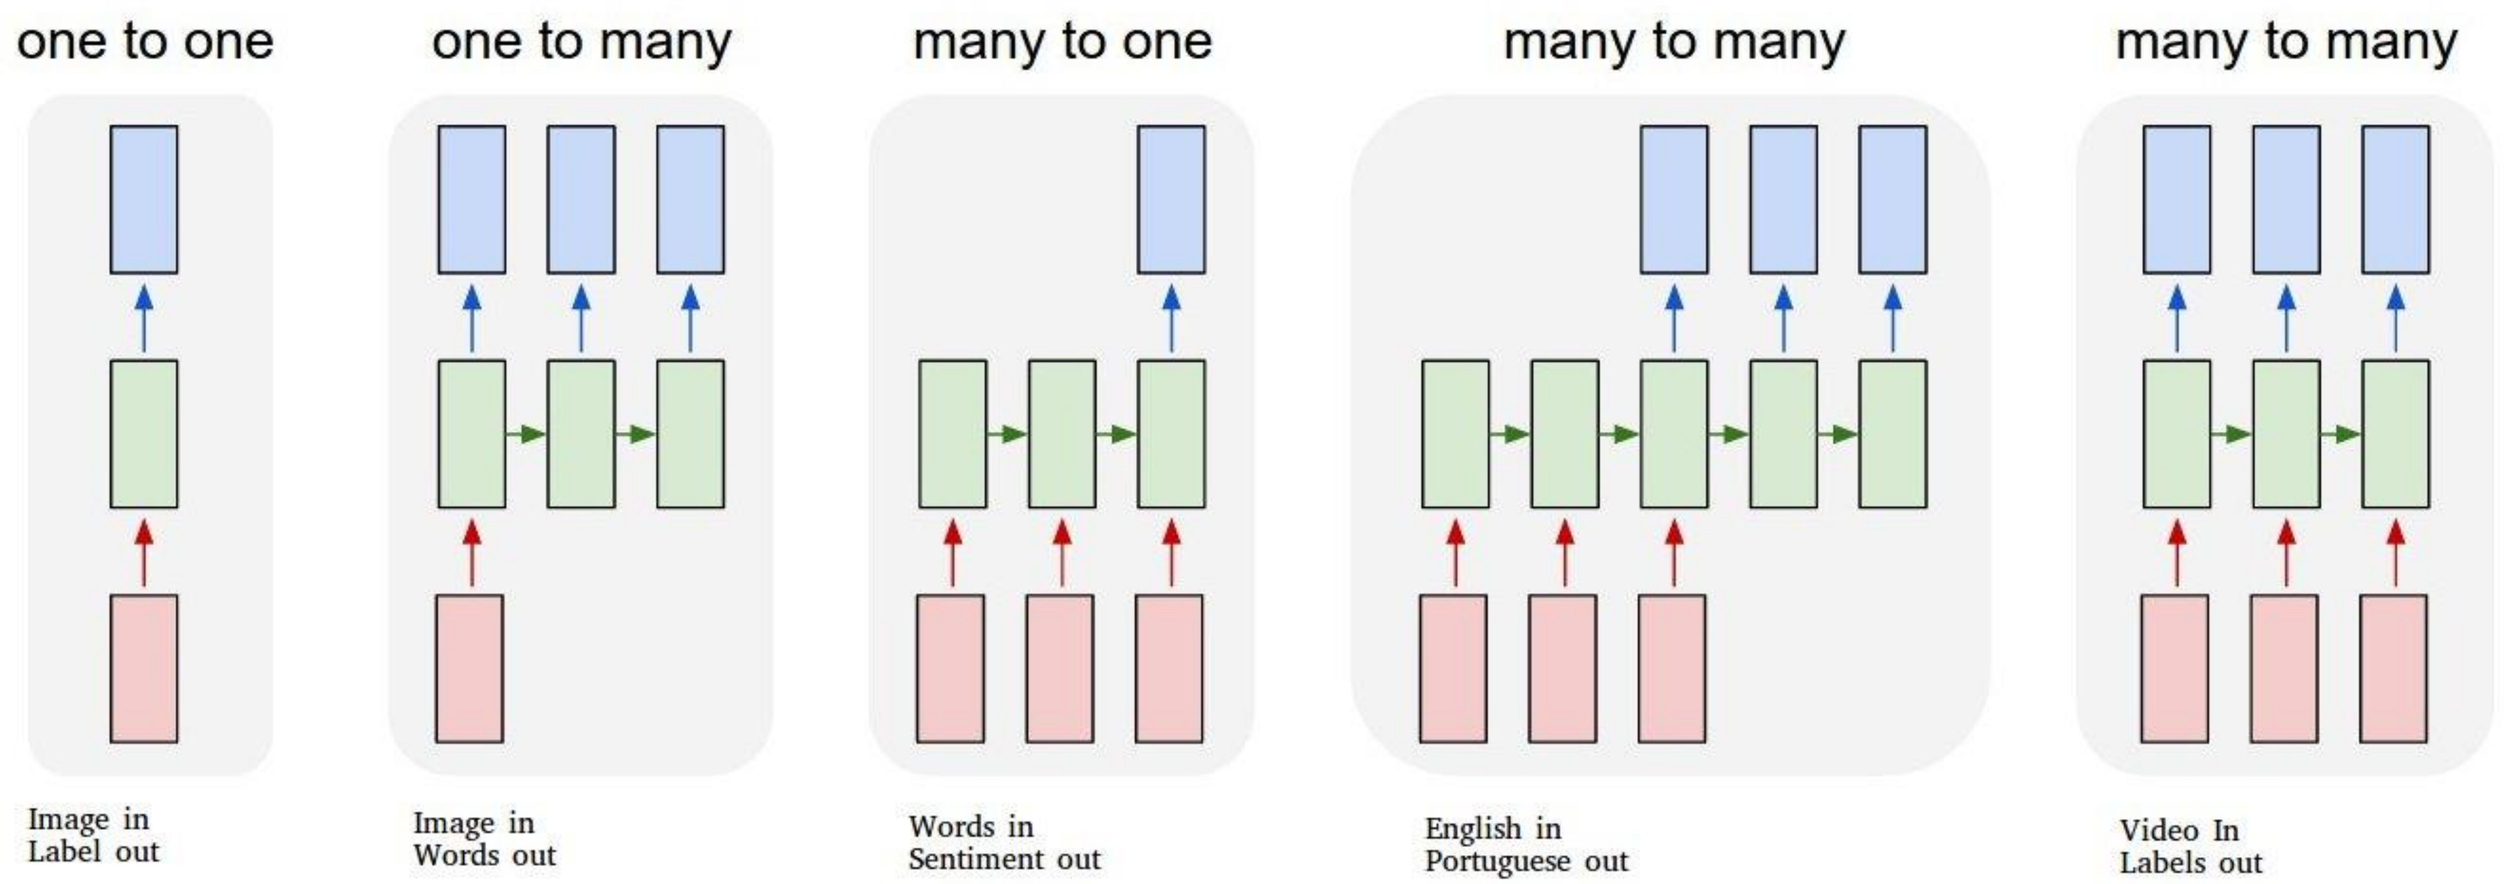

https://leonardoaraujosantos.gitbooks.io/artificial-inteligence/content/recurrent_neural_networks.html

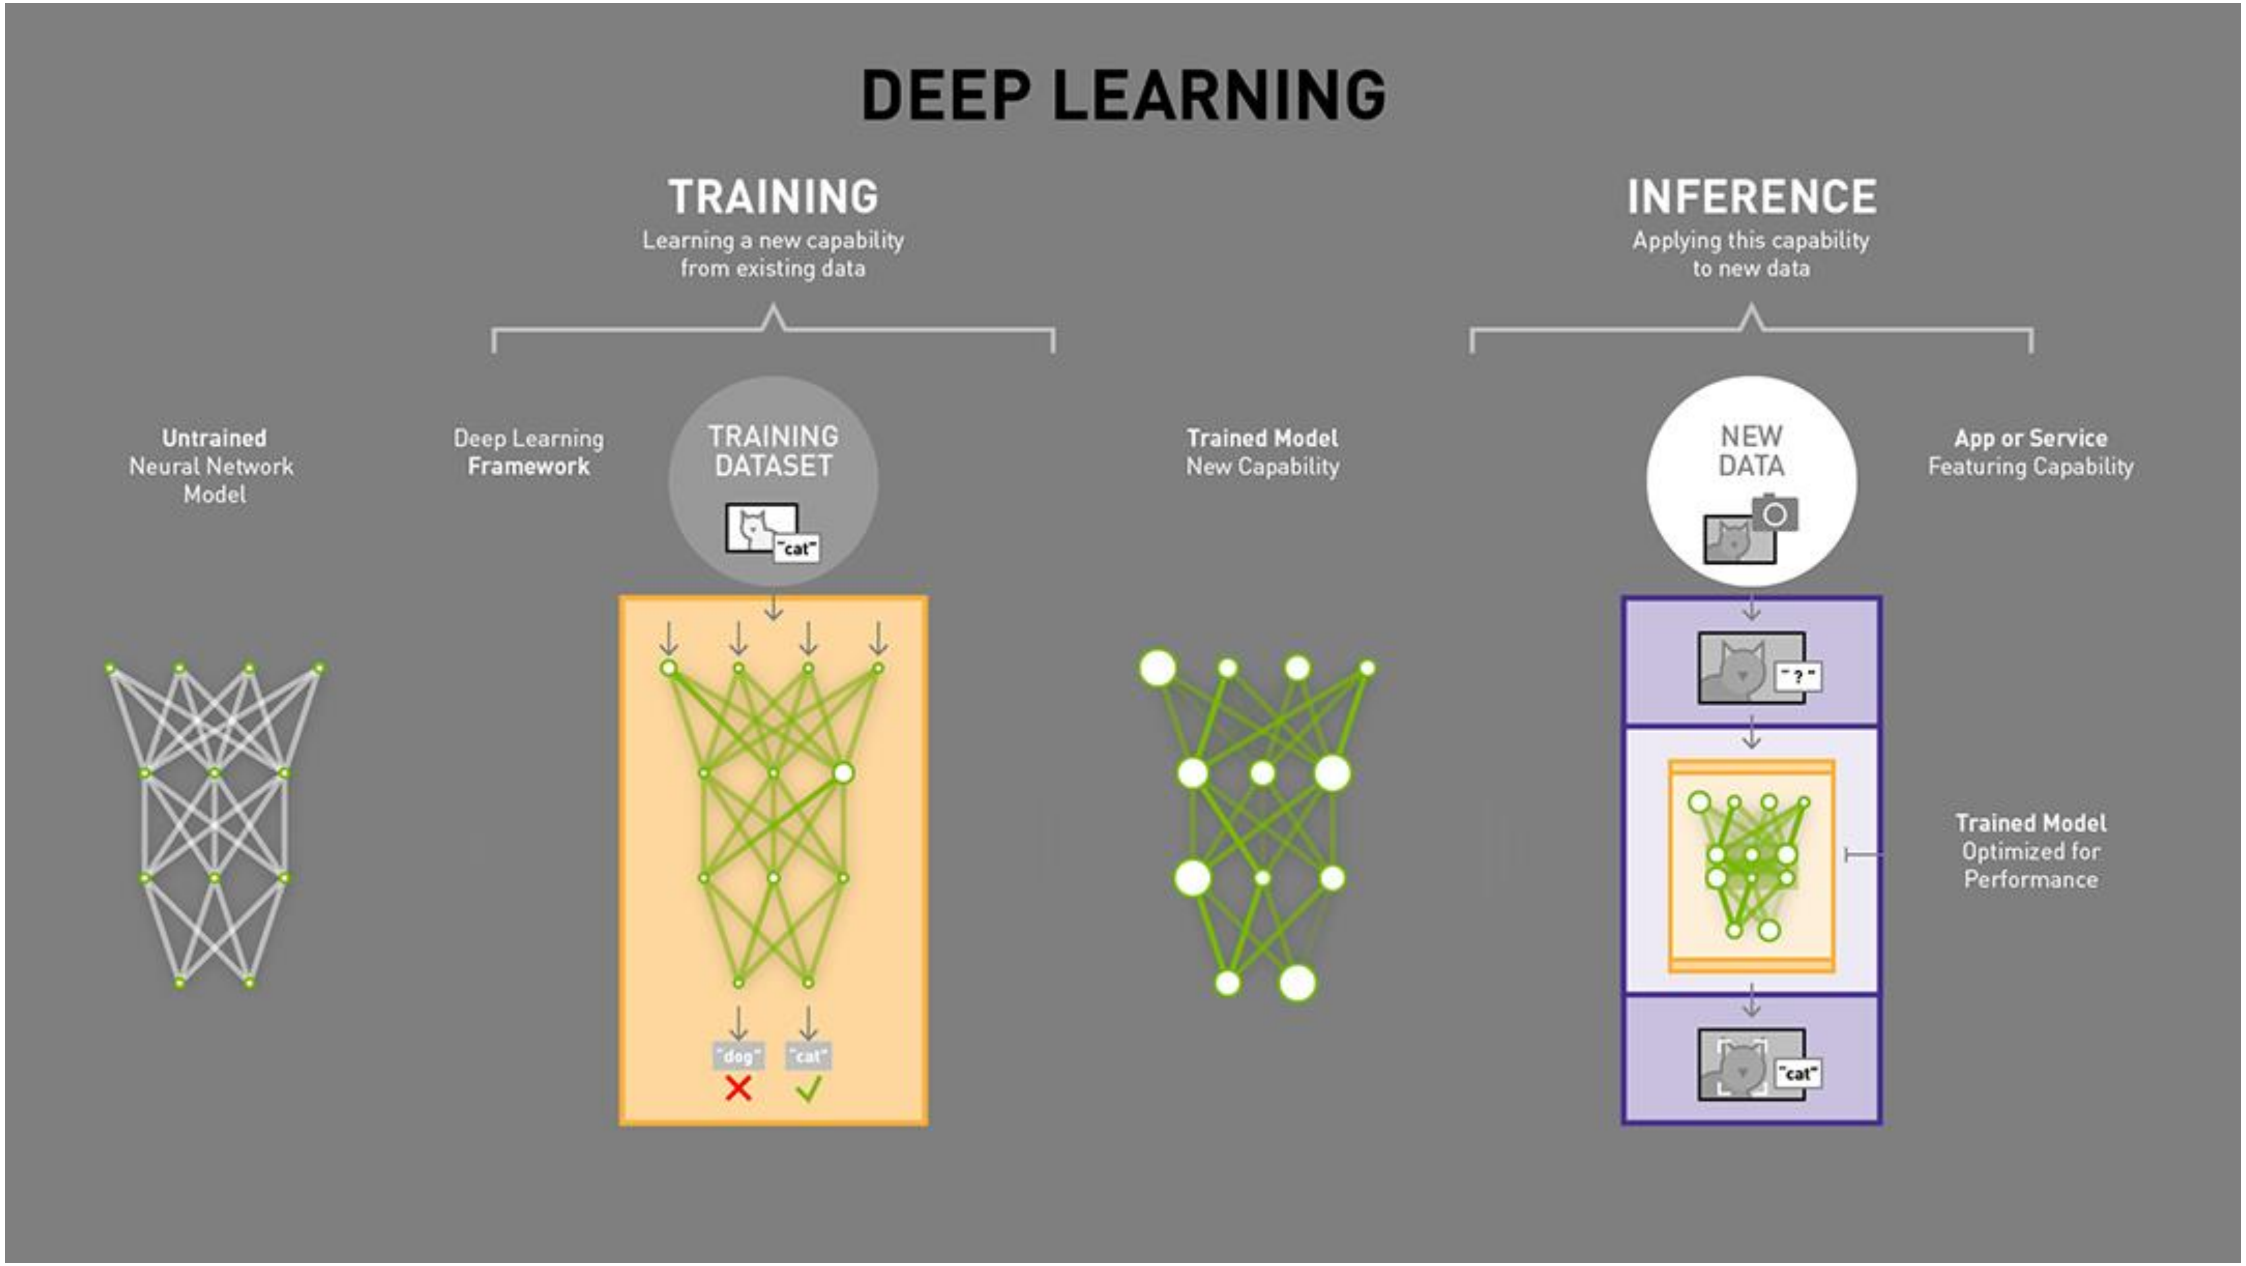

LINKÖPINGS UNIVERSITET

https://blogs.nvidia.com/blog/2016/08/22/difference-deep-learning-training-inference-ai/

Traditional vs ML problem solving

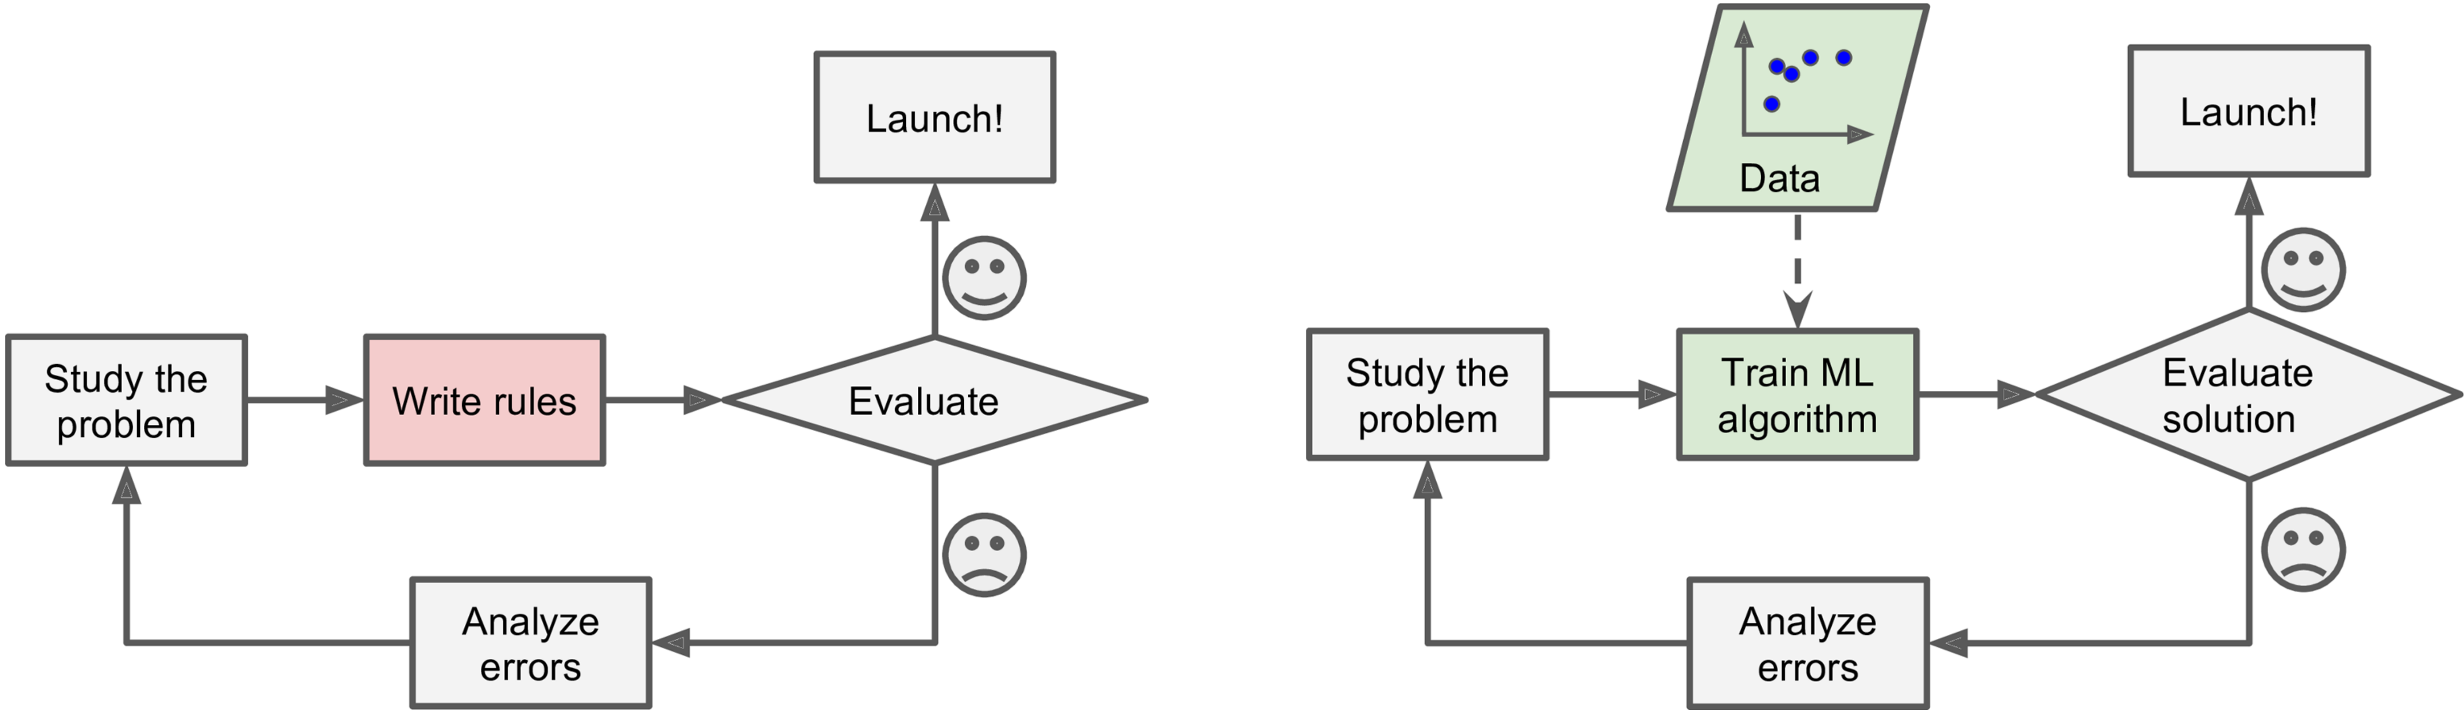

Image Classification

ImageNet Challenge

IMAGENET

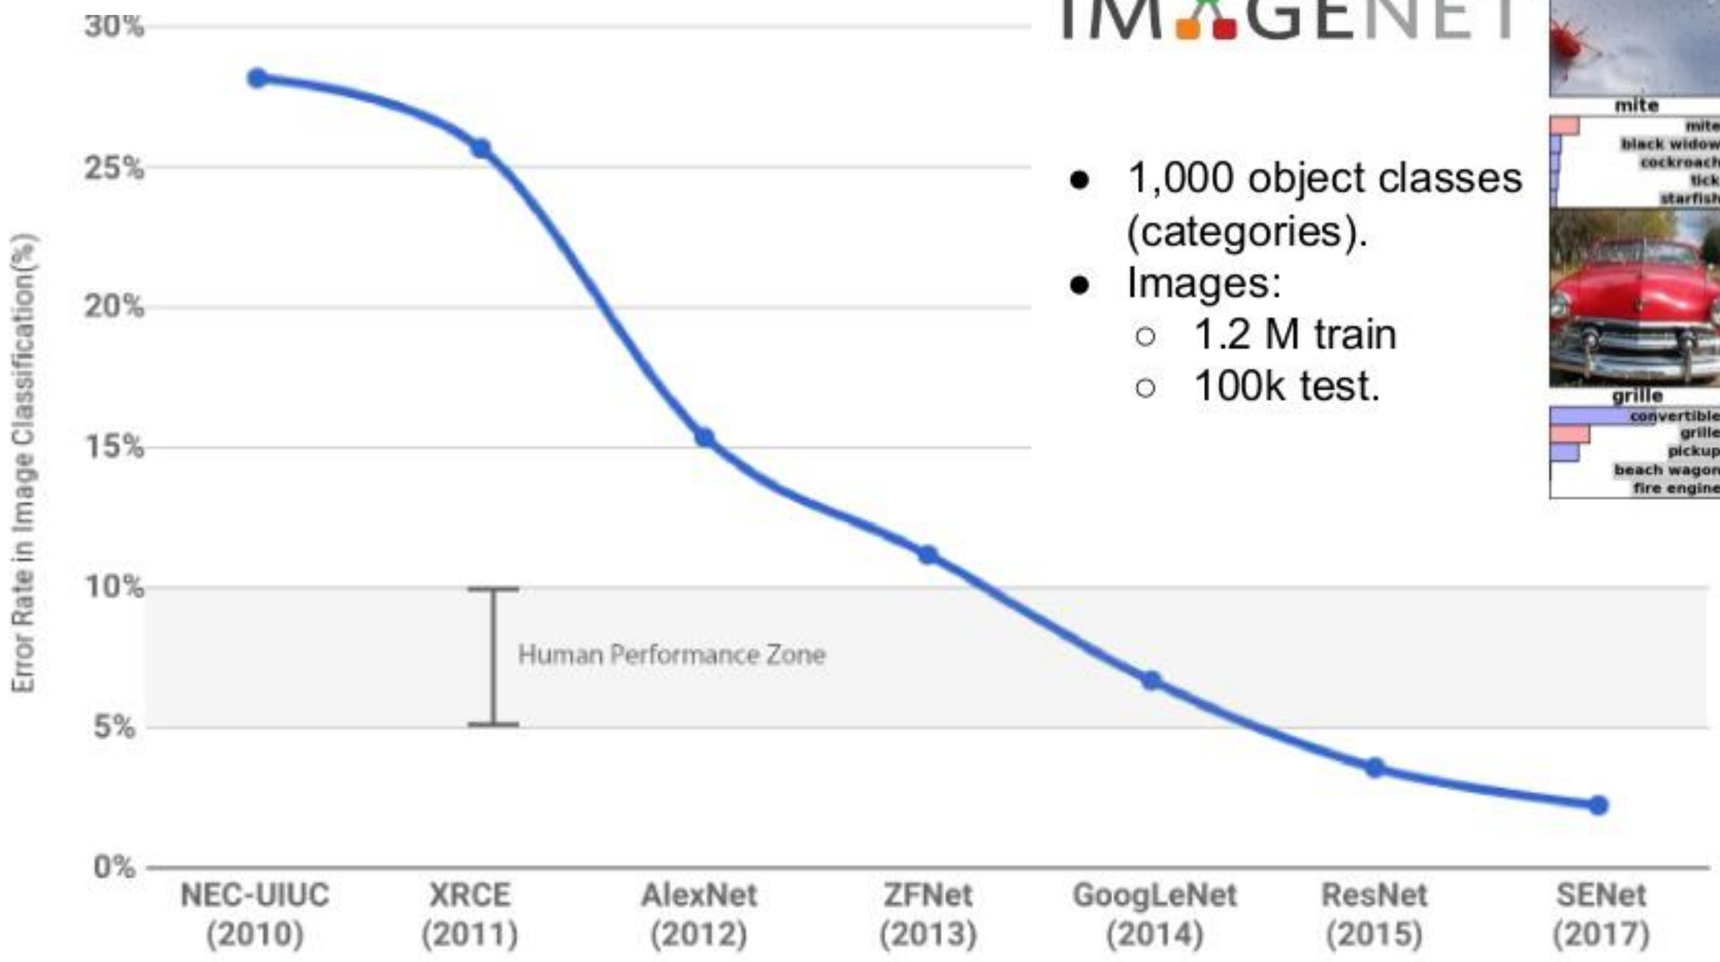

- 1,000 object classes (categories).
- Images:
  - 1.2 M train
  - 100k test.

Speech Recognition

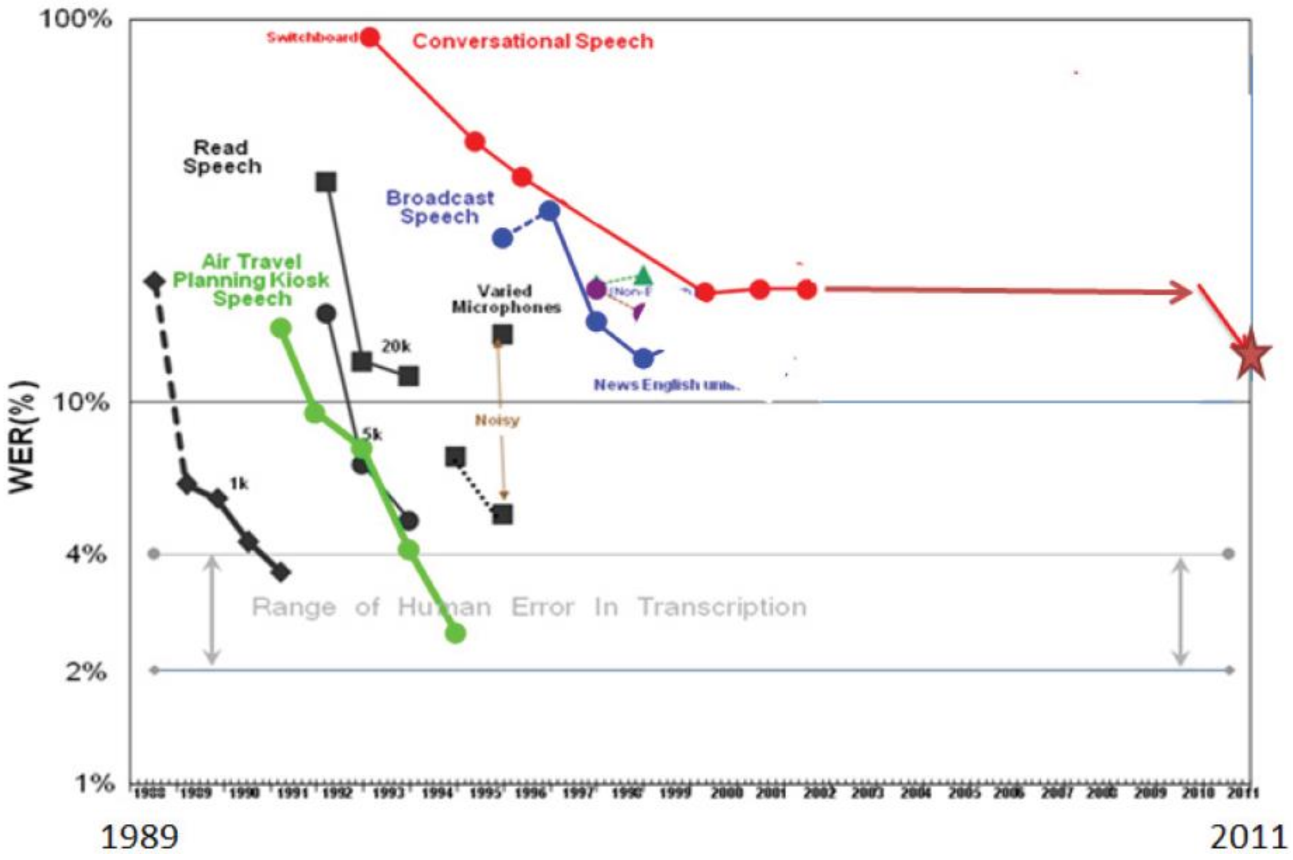

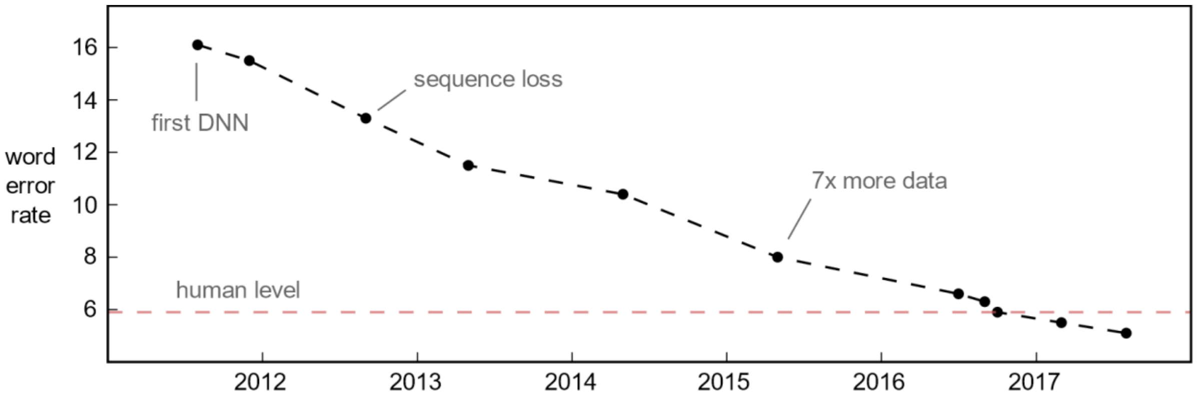

Neural Networks Timeline

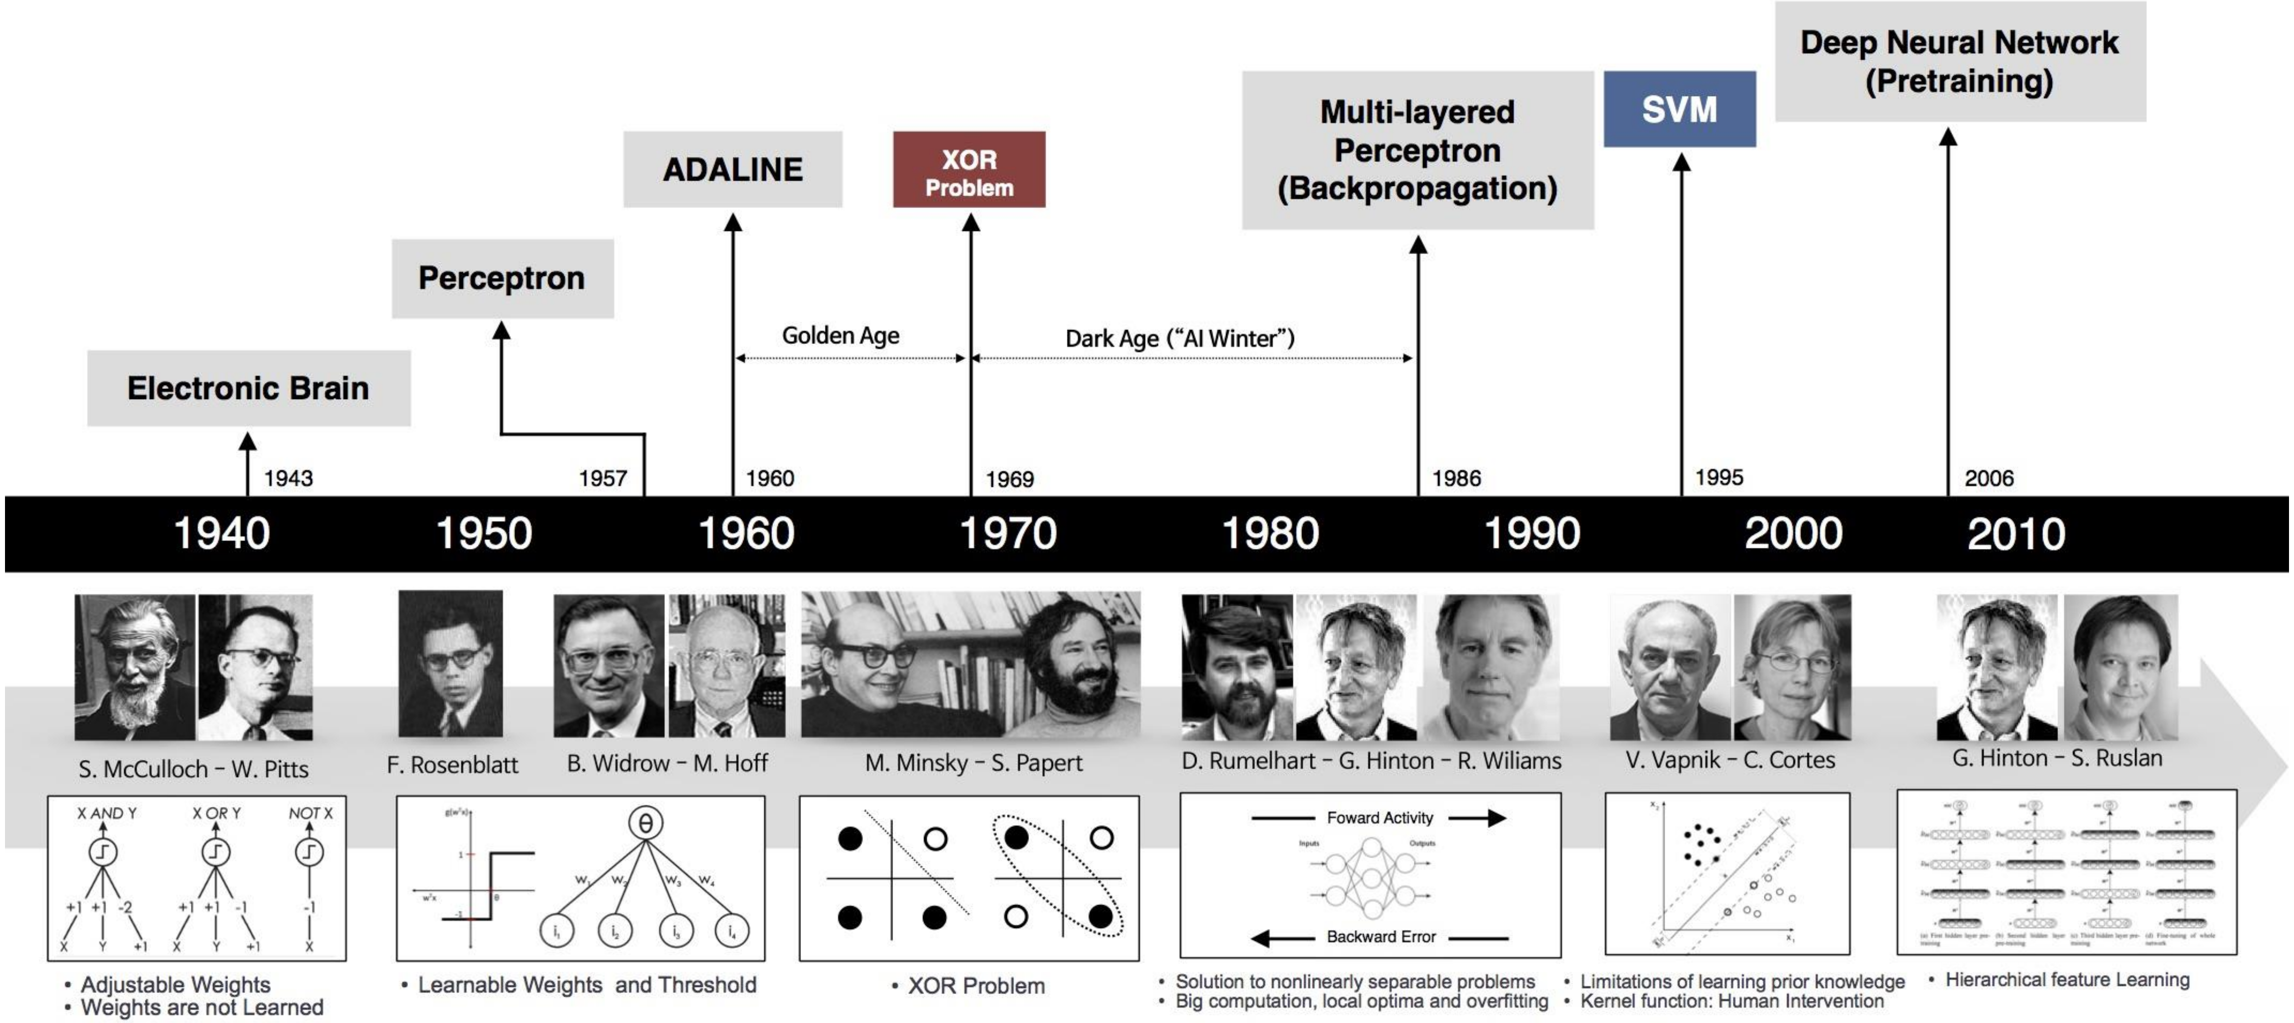

LINKÖPINGS UNIVERSITET

https://beamandrew.github.io/deeplearning/2017/02/23/deep_learning_101_part1.html

Training, Validation, and Test Data

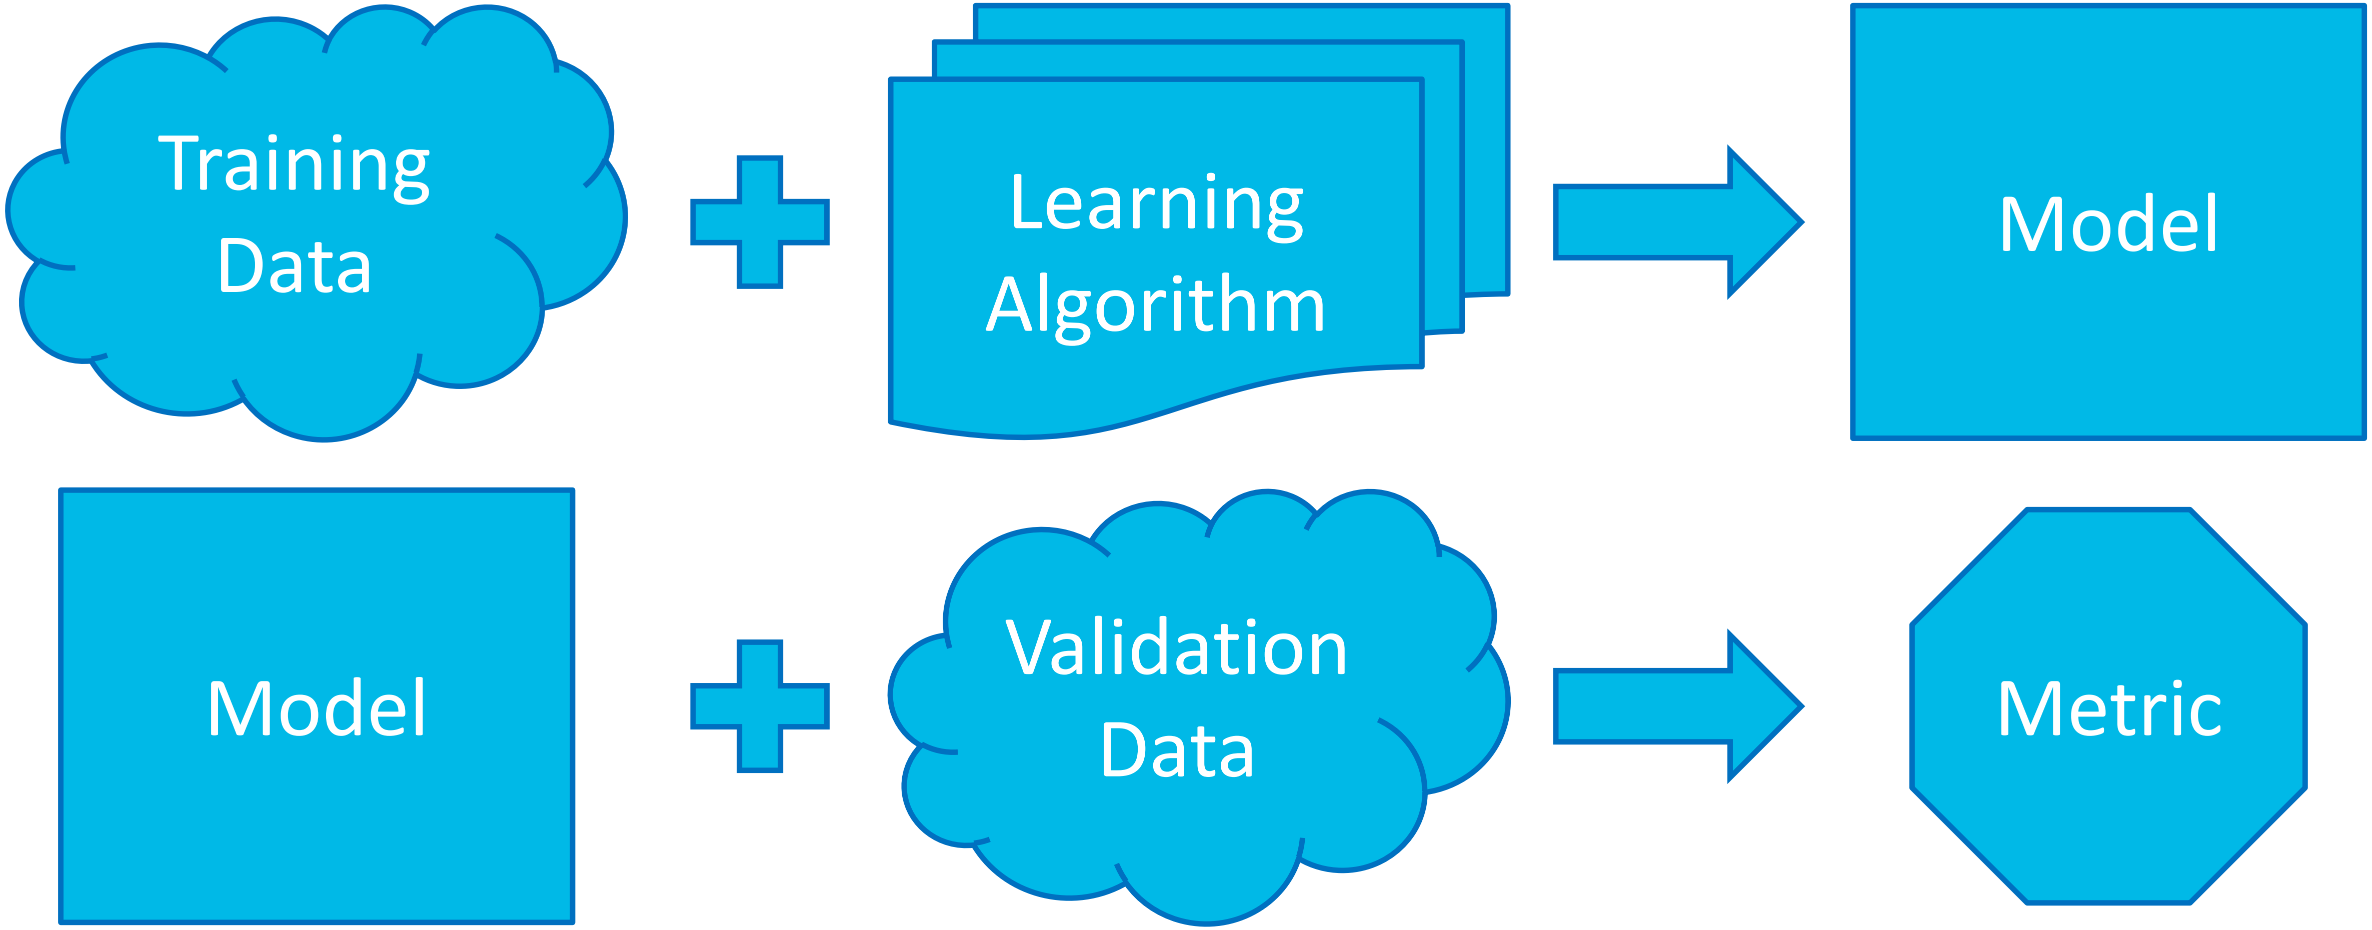

Precision and Recall

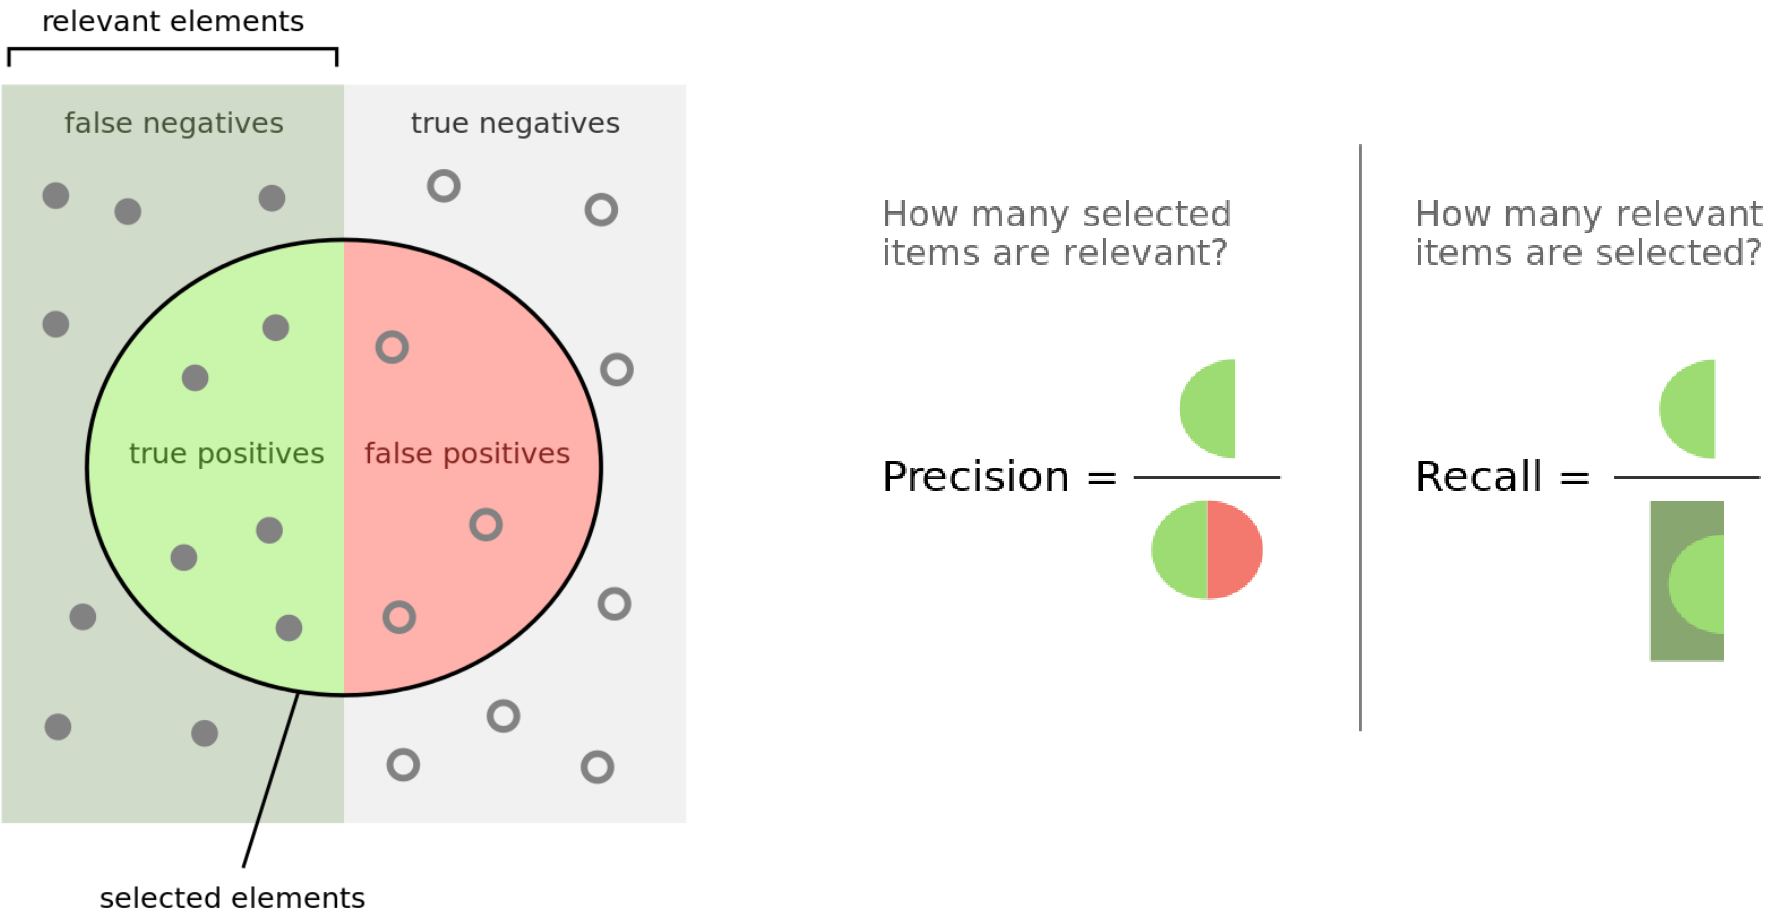

https://en.wikipedia.org/wiki/Precision_and_recall

Machine Learning Process

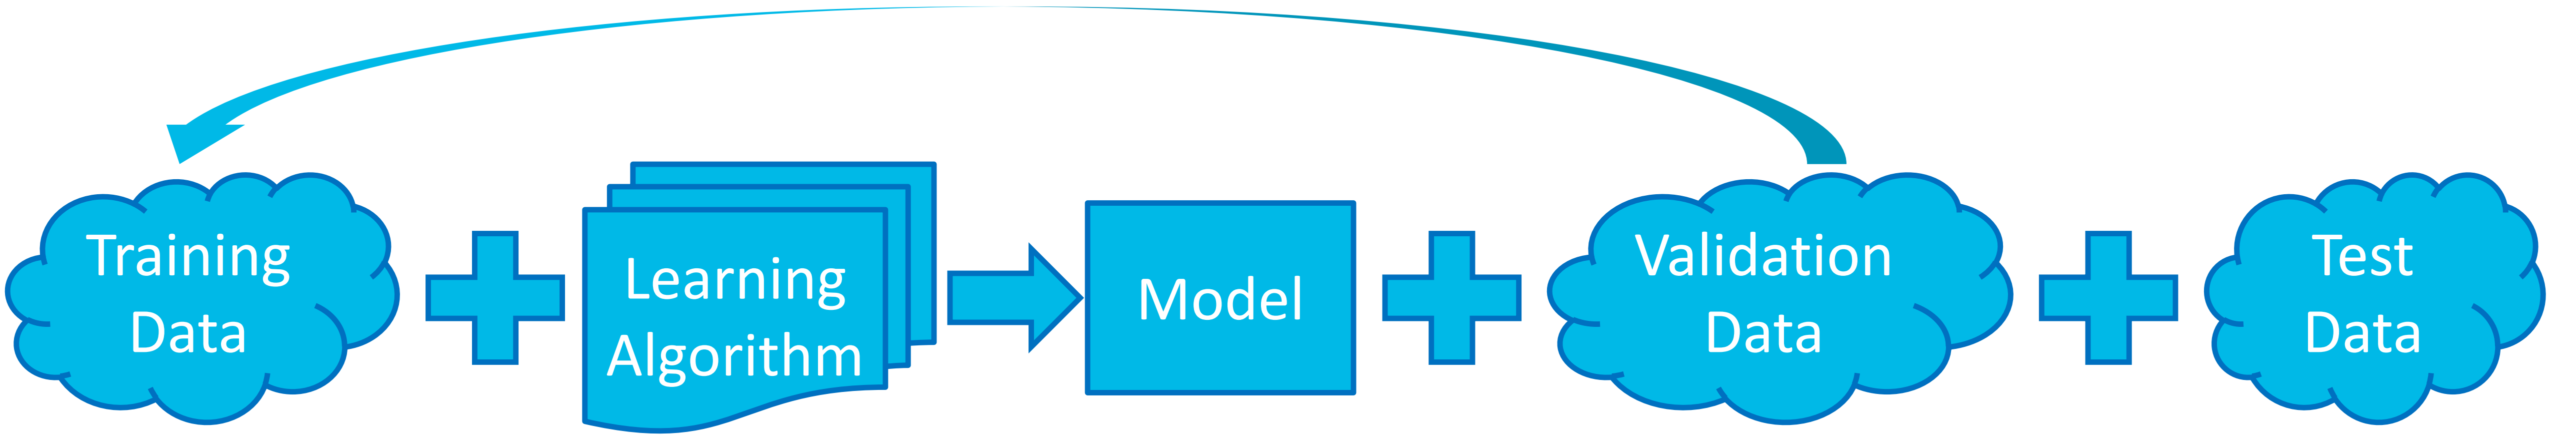

Underfitting and Overfitting

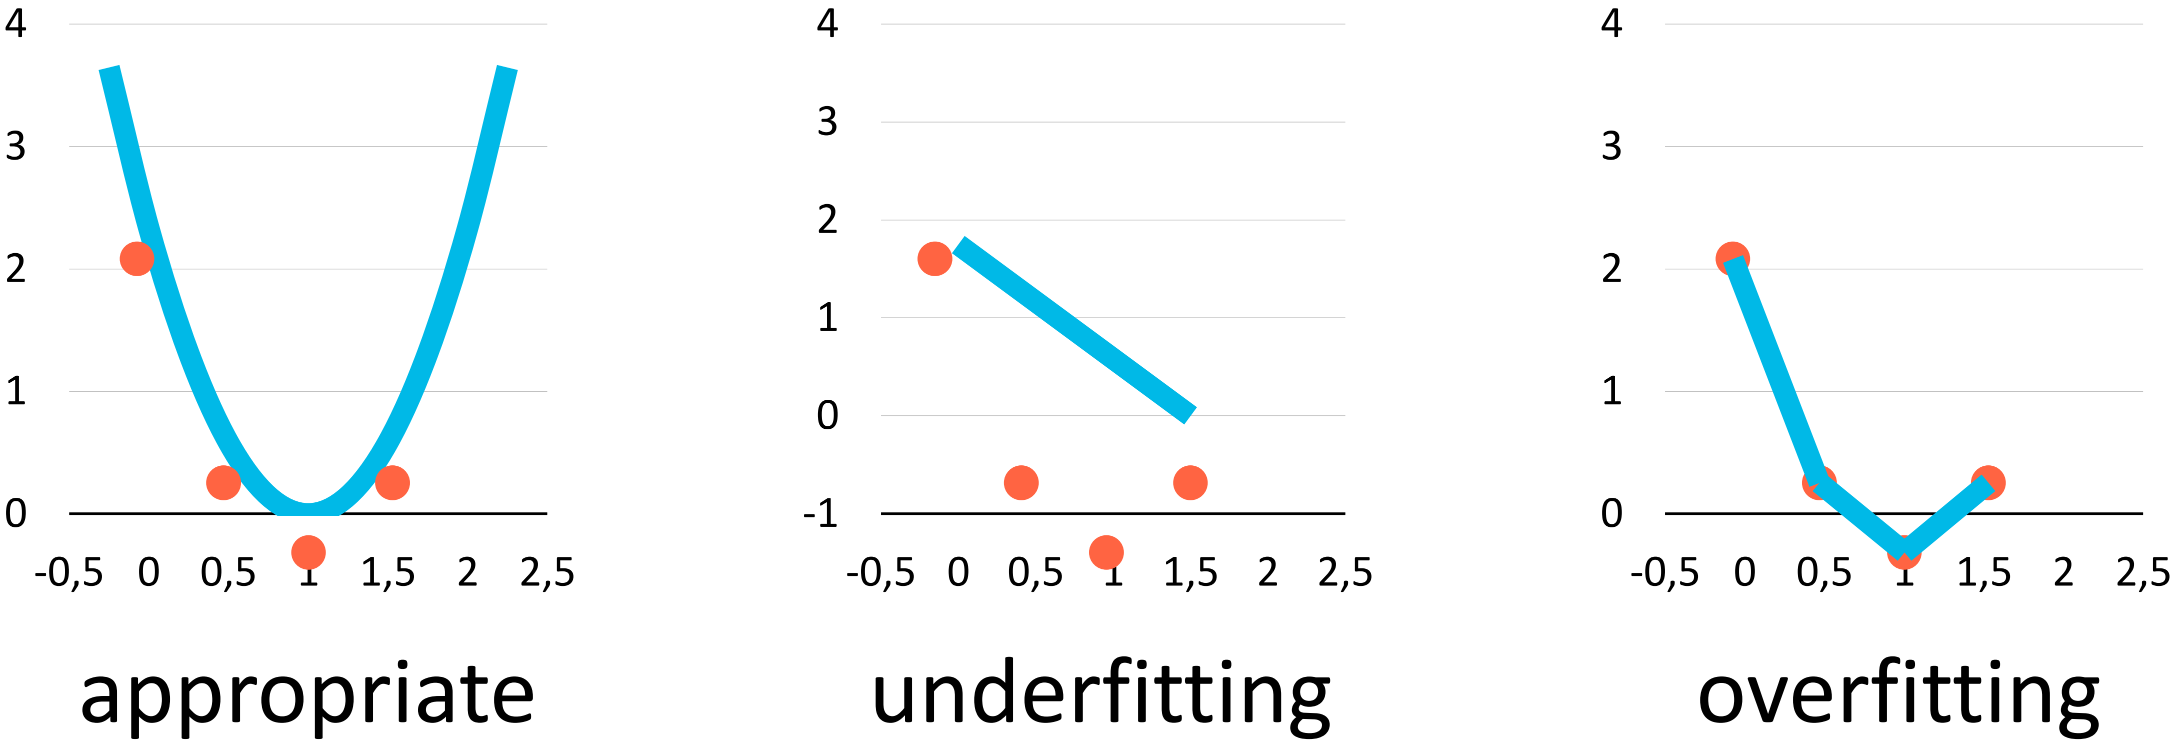

LINKÖPINGS UNIVERSITET

With a $p = n - 1$ degree polynomial, we can fit n data points perfectly.

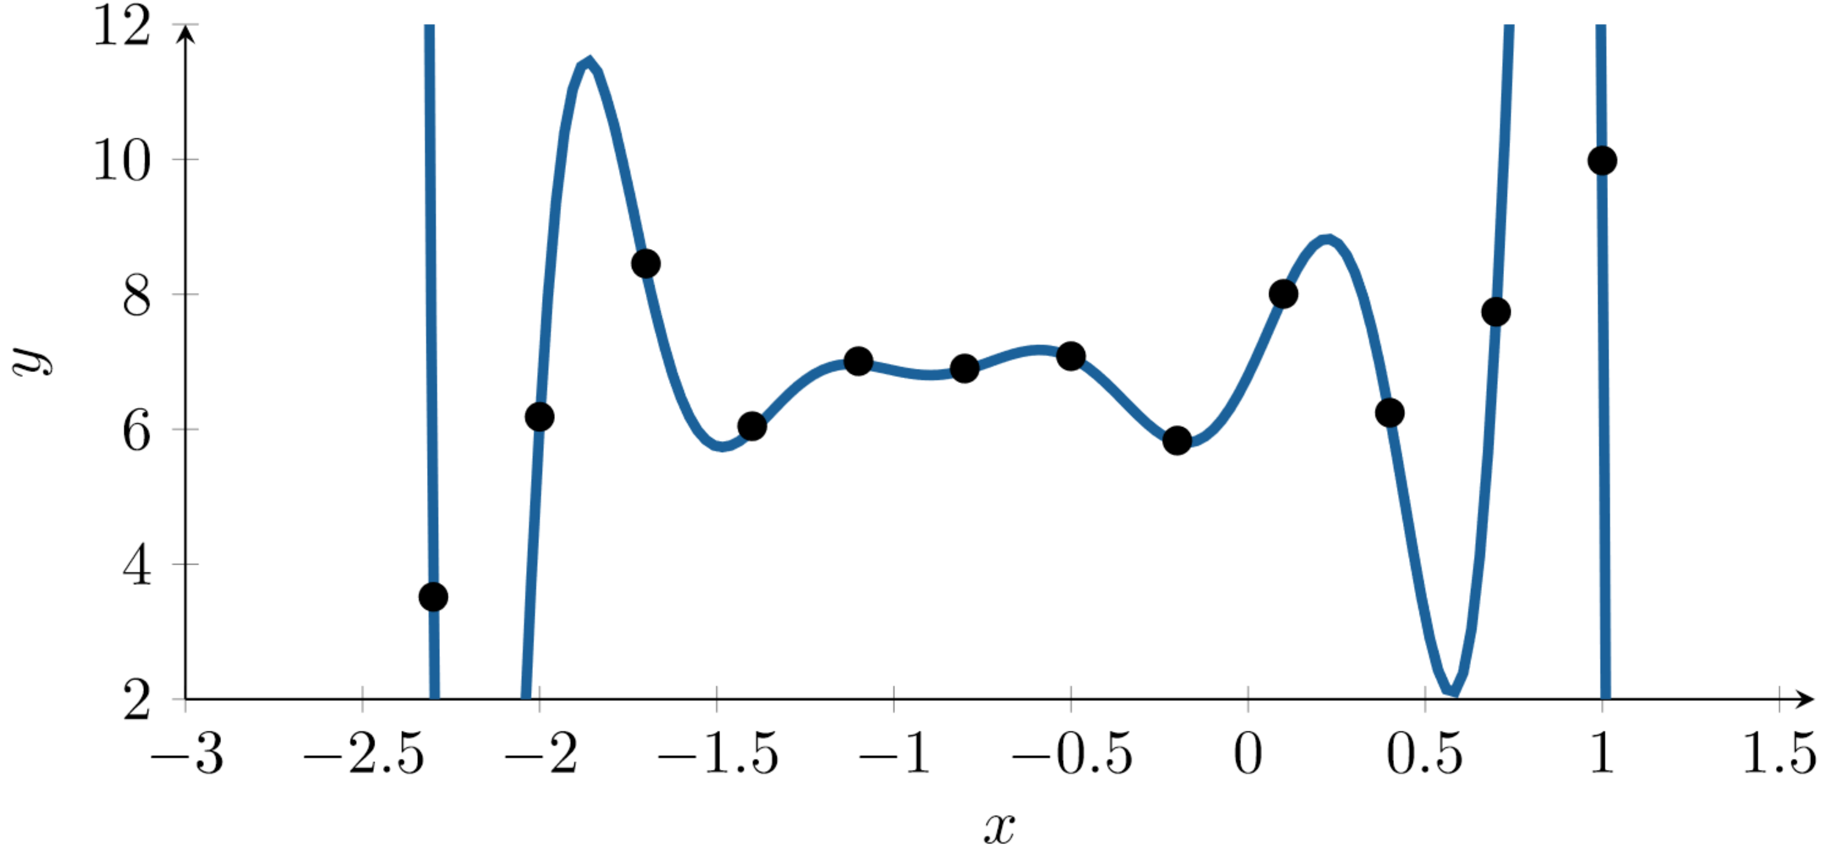

Regularization

"Keep β small unless the data really convinces us otherwise"

Least squares with Ridge regression

$$\widehat{\boldsymbol{\beta}} = \underset{\boldsymbol{\beta}}{\text{argmin}} \|\mathbf{X}\boldsymbol{\beta} - \mathbf{y}\|_2^2 + \gamma\|\boldsymbol{\beta}\|_2^2$$

$$\Rightarrow (\mathbf{X}^{\mathsf{T}}\mathbf{X} + \gamma\mathbf{I}_{p+1})\widehat{\boldsymbol{\beta}} = \mathbf{X}^{\mathsf{T}}\mathbf{y}$$

γ regularization parameter

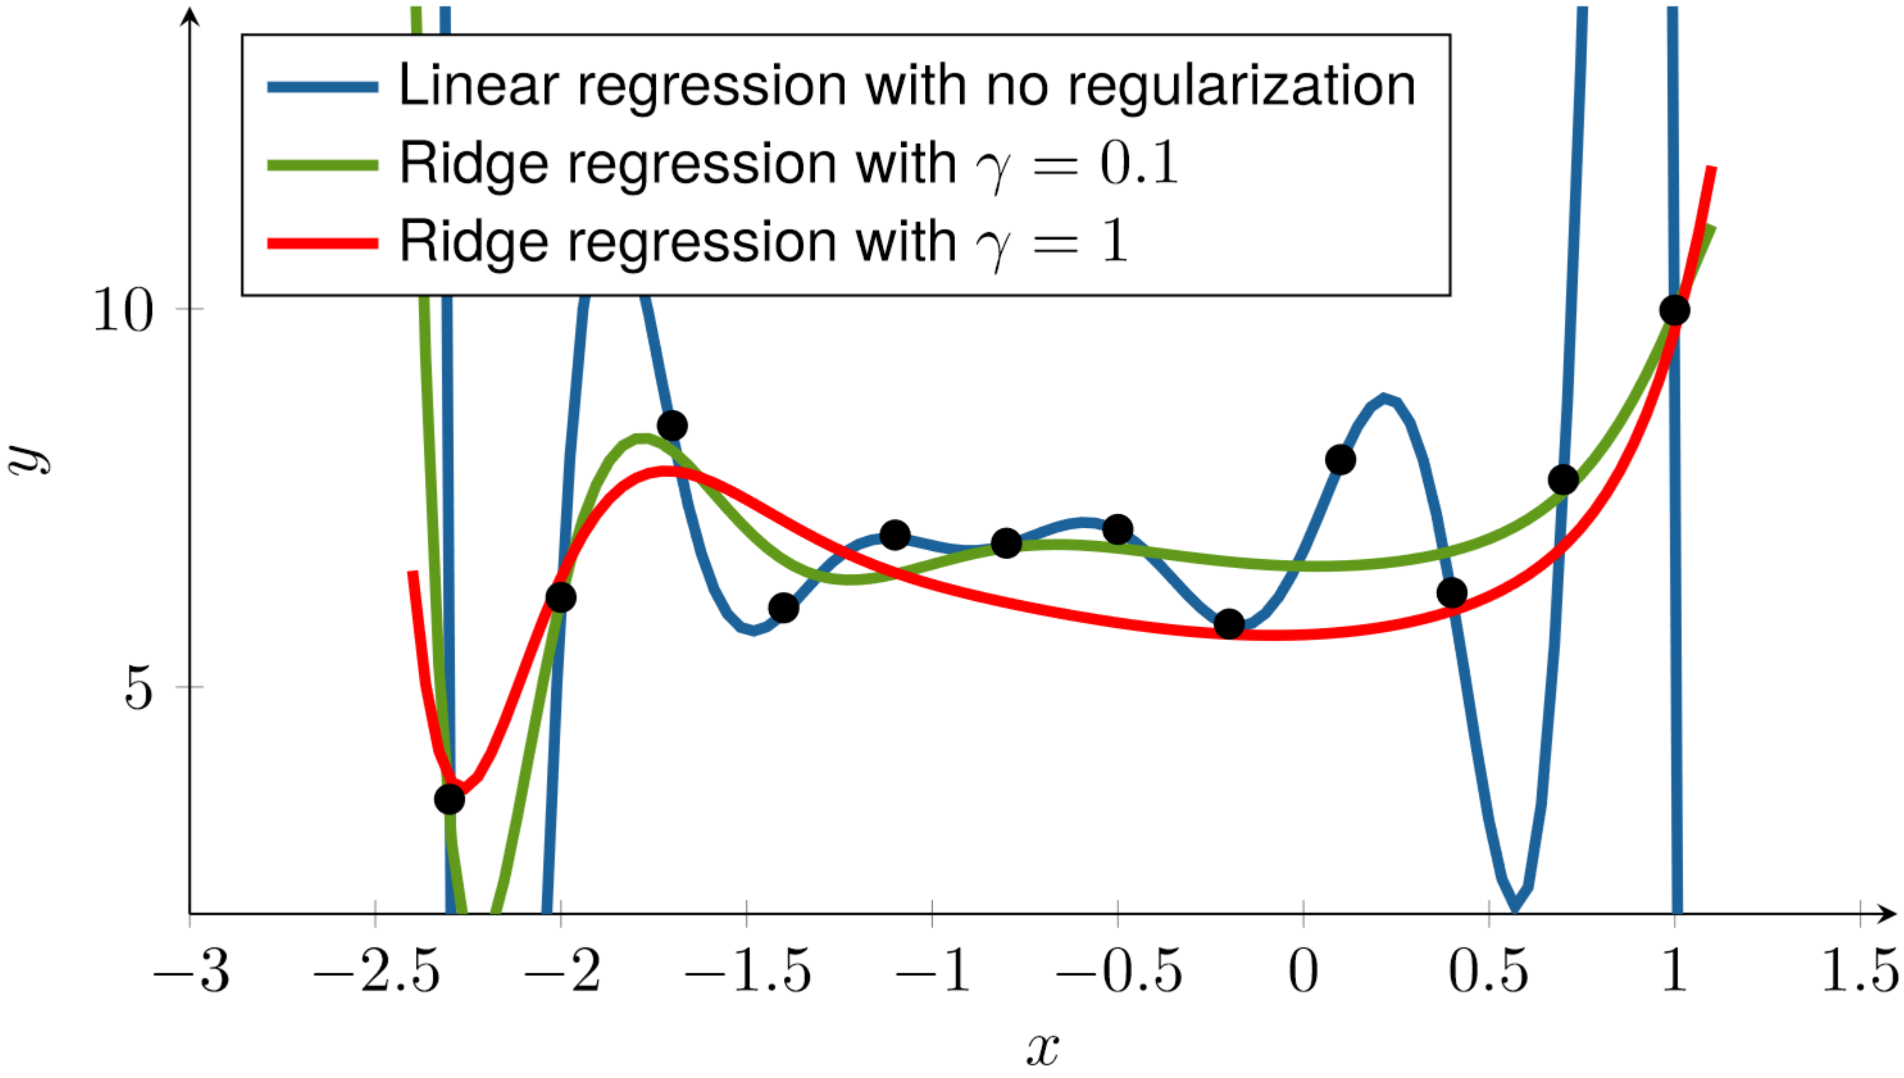

Regularization can help us to avoid overfitting!

Relationship between model capacity and error

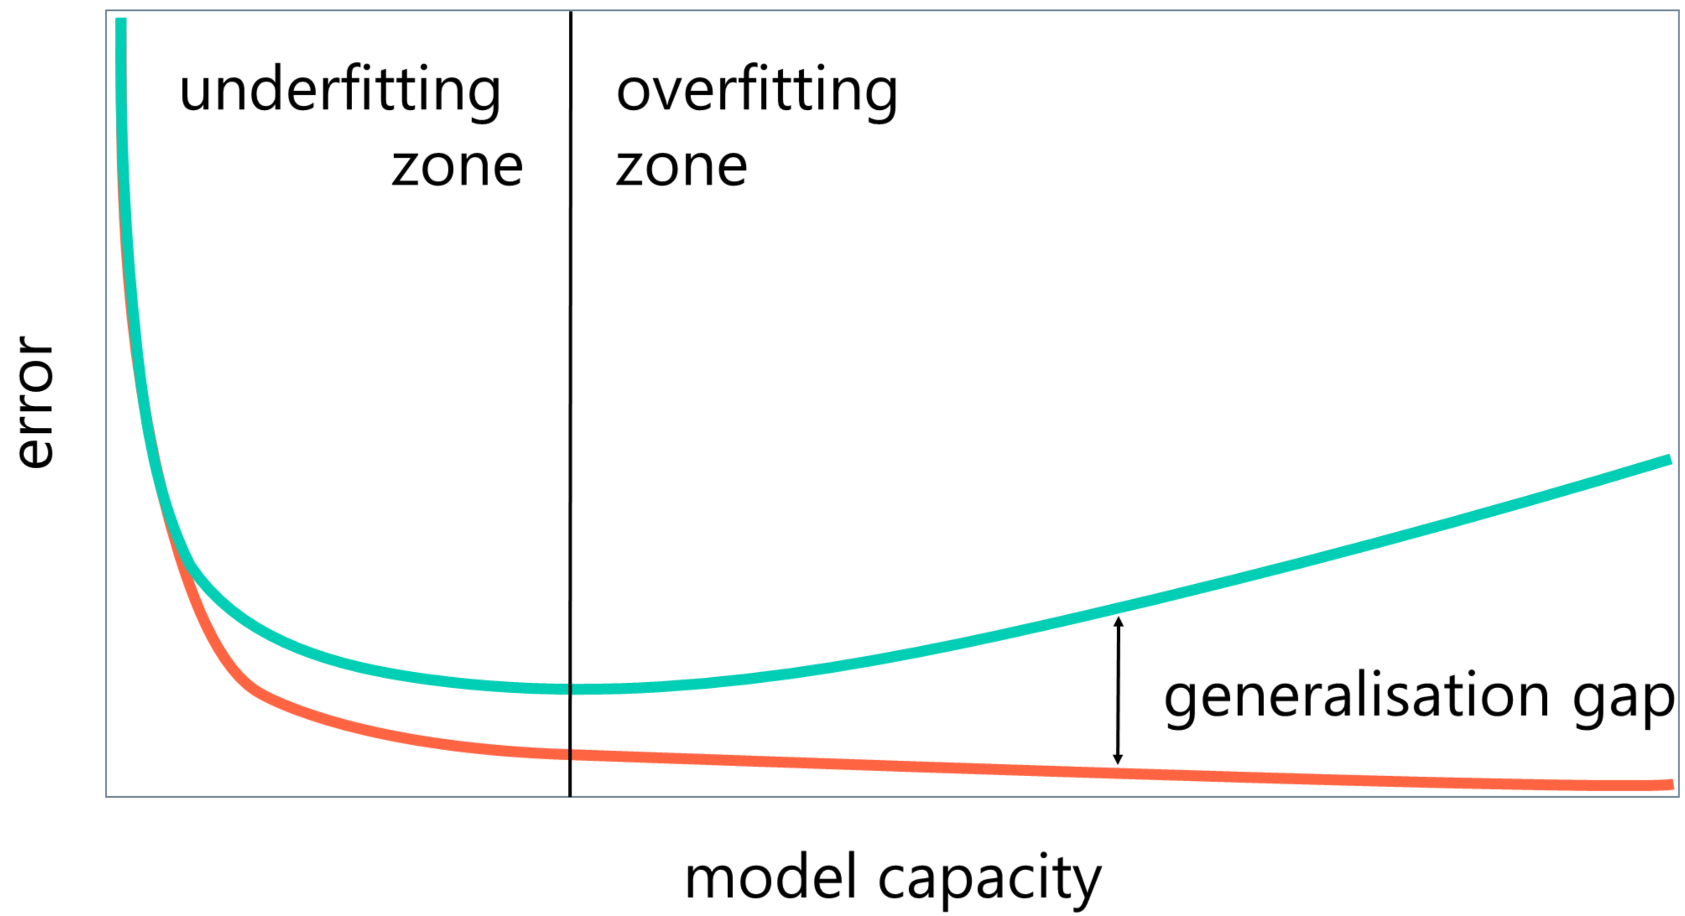

Bias

- **Dataset bias** – When the data used to train machine learning models doesn't represent the diversity of the customer base.
- **Association bias** – When the data used to train a model reinforces and multiplies a cultural bias.
- **Automation bias** – When automated decisions override social and cultural considerations.
- **Interaction bias** – When humans tamper with AI and create biased results.
- **Confirmation bias** – When oversimplified personalization makes biased assumptions for a group or an individual.

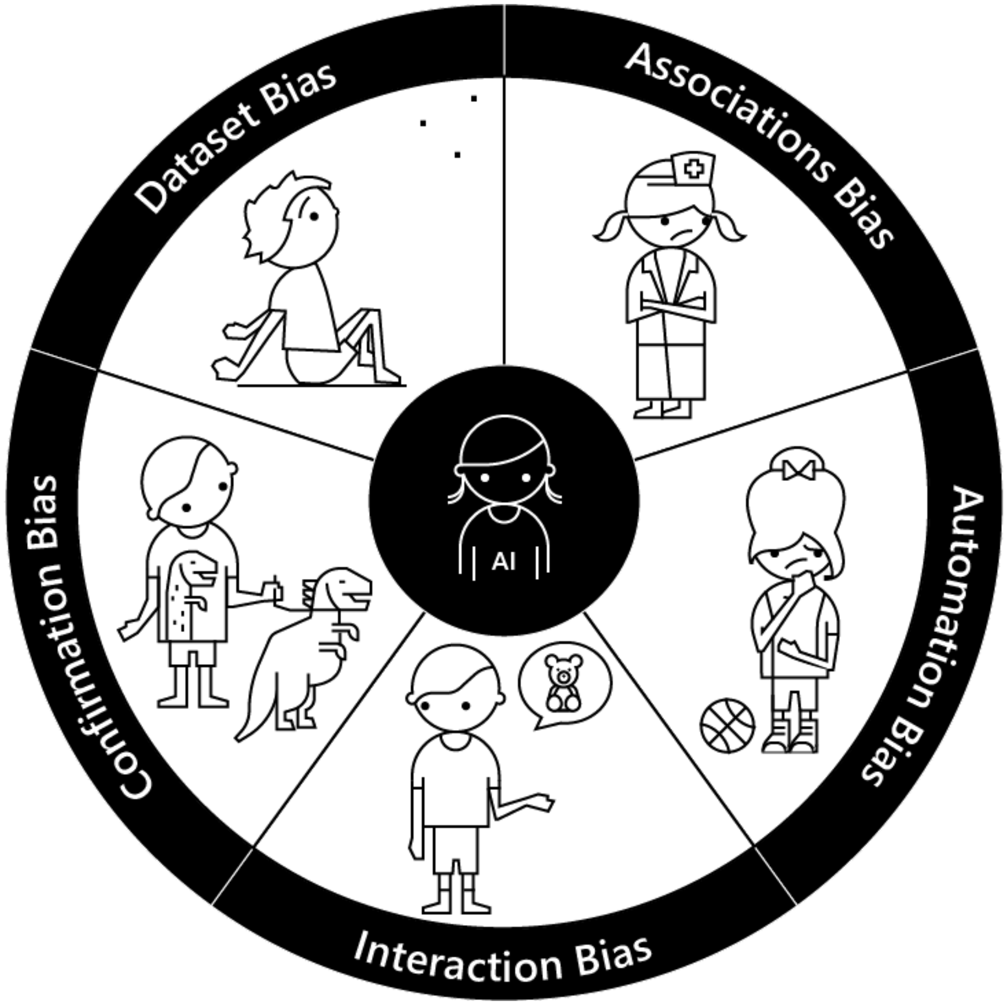

https://medium.com/microsoft-design/how-to-recognize-exclusion-in-ai-ec2d6d89f850

Machine learning is still brittle…

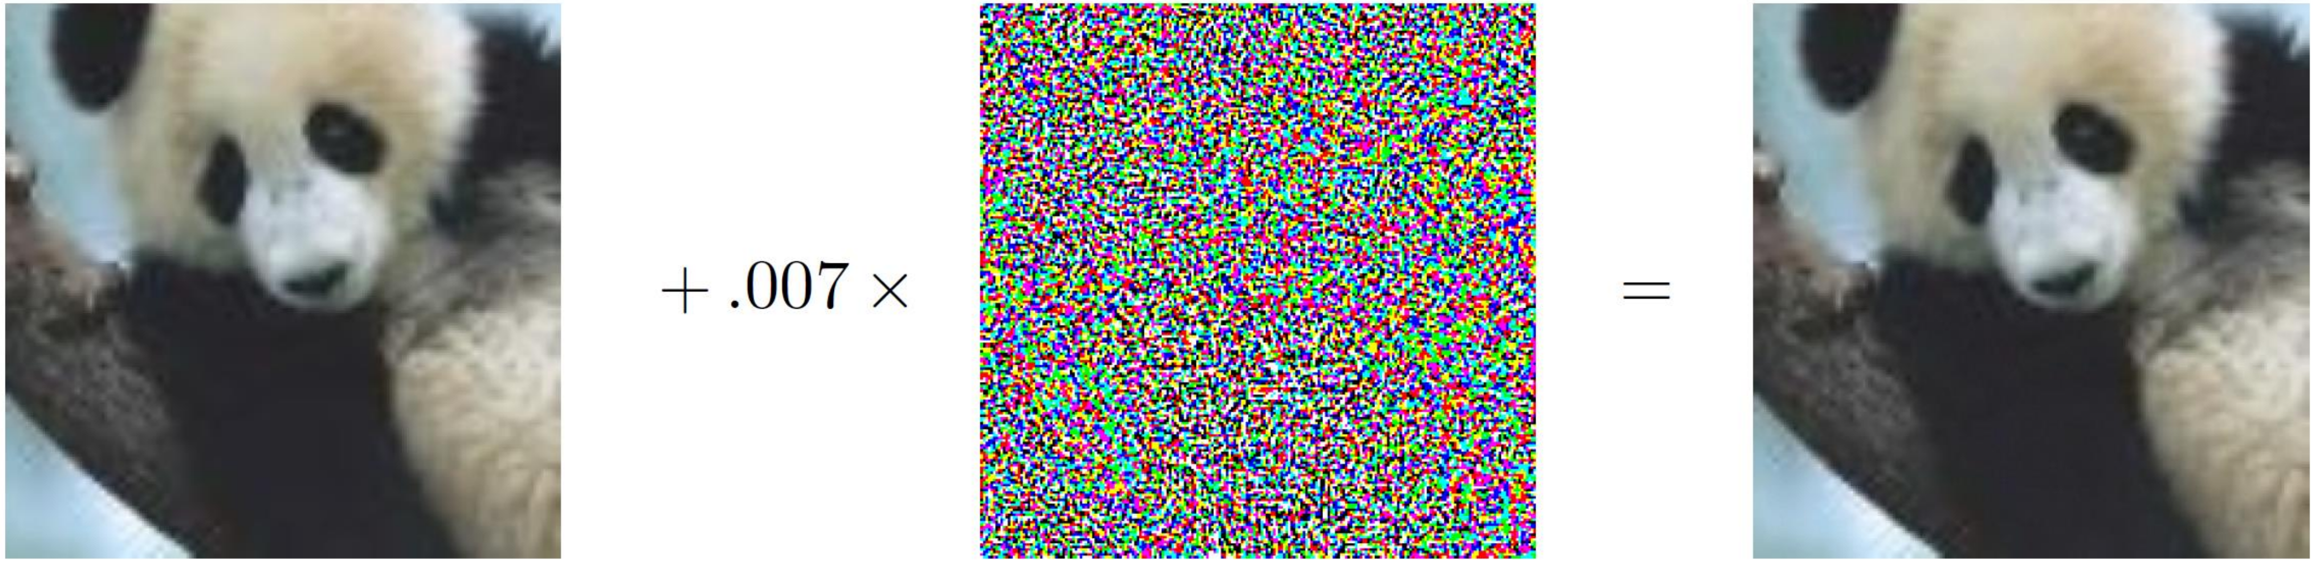

$\boldsymbol{x}$
"panda"
57.7% confidence

$\text{sign}(\nabla_{\boldsymbol{x}} J(\boldsymbol{\theta}, \boldsymbol{x}, y))$
"nematode"
8.2% confidence

$\boldsymbol{x} + \epsilon\text{sign}(\nabla_{\boldsymbol{x}} J(\boldsymbol{\theta}, \boldsymbol{x}, y))$
"gibbon"
99.3 % confidence

Ian J. Goodfellow, Jonathon Shlens, Christian Szegedy. Explaining and Harnessing Adversarial Examples. ICLR 2015
https://arxiv.org/abs/1412.6572

Generative Adversarial Networks (GANs)

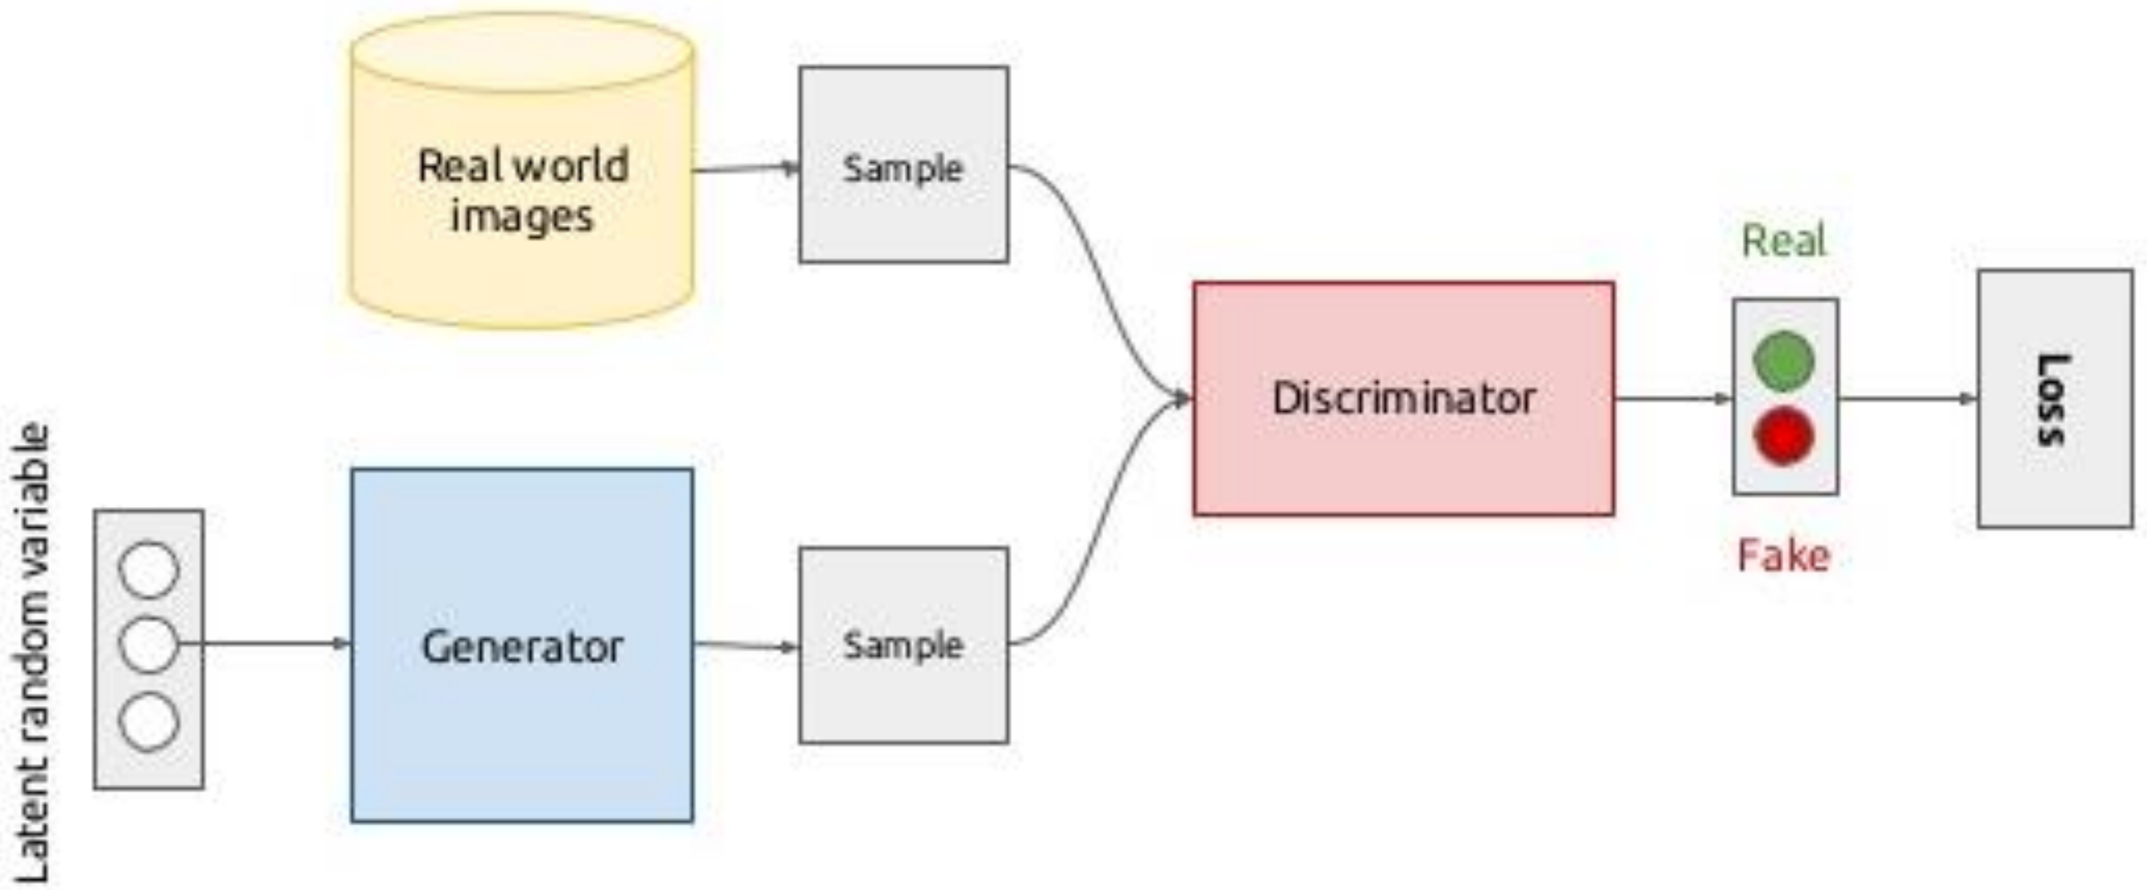

Kevin McGuinness. Deep Learning for Computer Vision: Generative models and adversarial training (UPC 2016). http://www.slideshare.net/xavigiro/deep-learning-for-computer-vision-generative-models-and-adversarial-training-upc-2016

LINKÖPINGS UNIVERSITET

https://www.thispersondoesnotexist.com/

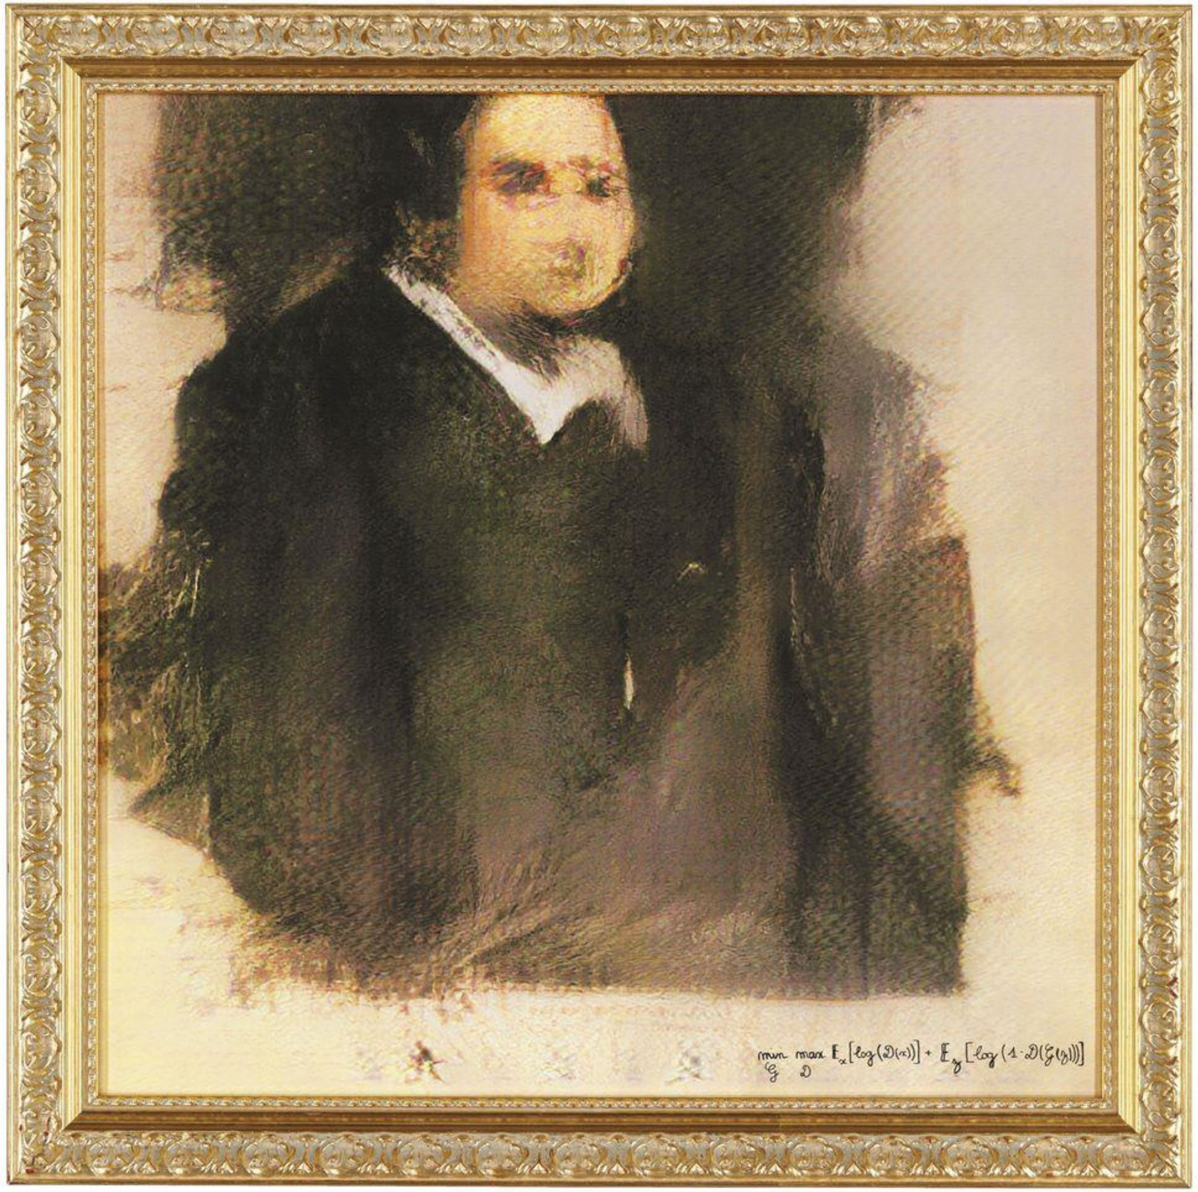

AI-Generated Portrait Sells for $432,500

https://www.bloomberg.com/news/articles/2018-10-25/ai-generated-portrait-is-sold-for-432-500-in-an-auction-first

Right for the Wrong Reasons [McCoy, Pavlick, Linzen ACL 2019]

Heuristic	Definition	Example
Lexical overlap	Assume that a premise entails all hypotheses constructed from words in the premise	**The doctor** was **paid** by **the actor**. $\xrightarrow[\text{WRONG}]{}$ The doctor paid the actor.
Subsequence	Assume that a premise entails all of its contiguous subsequences.	The doctor near **the actor danced**. $\xrightarrow[\text{WRONG}]{}$ The actor danced.
Constituent	Assume that a premise entails all complete subtrees in its parse tree.	If **the artist slept**, the actor ran. $\xrightarrow[\text{WRONG}]{}$ The artist slept.

Table 1: The heuristics targeted by the HANS dataset, along with examples of incorrect entailment predictions that these heuristics would lead to.

Right for the Wrong Reasons: Diagnosing Syntactic Heuristics in Natural Language Inference
by R. Thomas McCoy, Ellie Pavlick, Tal Linzen https://arxiv.org/abs/1902.01007

Right for the Wrong Reasons [McCoy, Pavlick, Linzen ACL 2019]

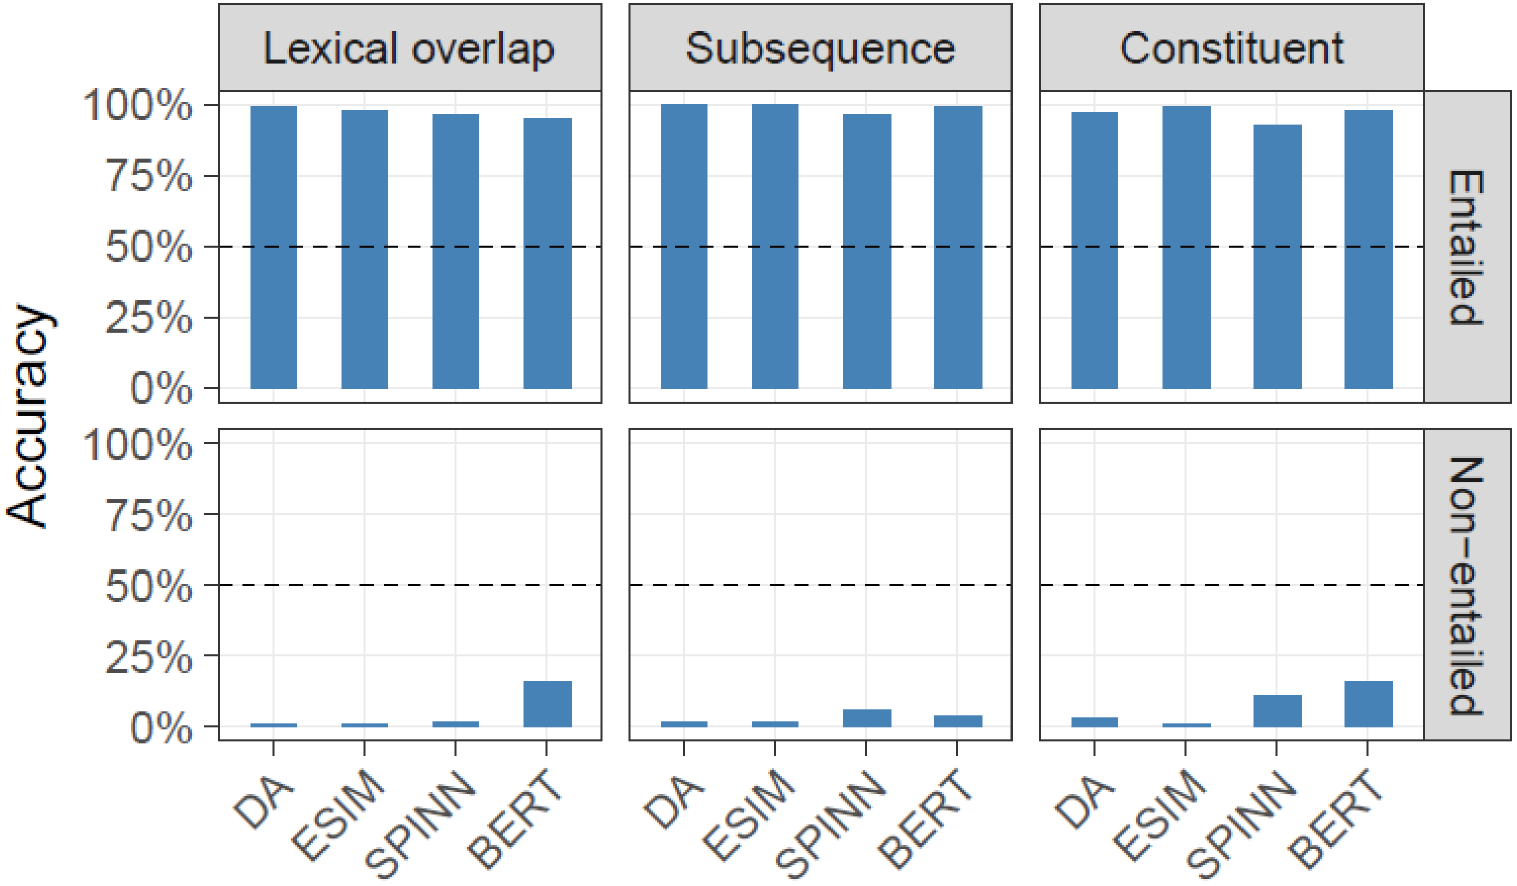

Right for the Wrong Reasons: Diagnosing Syntactic Heuristics in Natural Language Inference
by R. Thomas McCoy, Ellie Pavlick, Tal Linzen https://arxiv.org/abs/1902.01007

The bigger system / picture

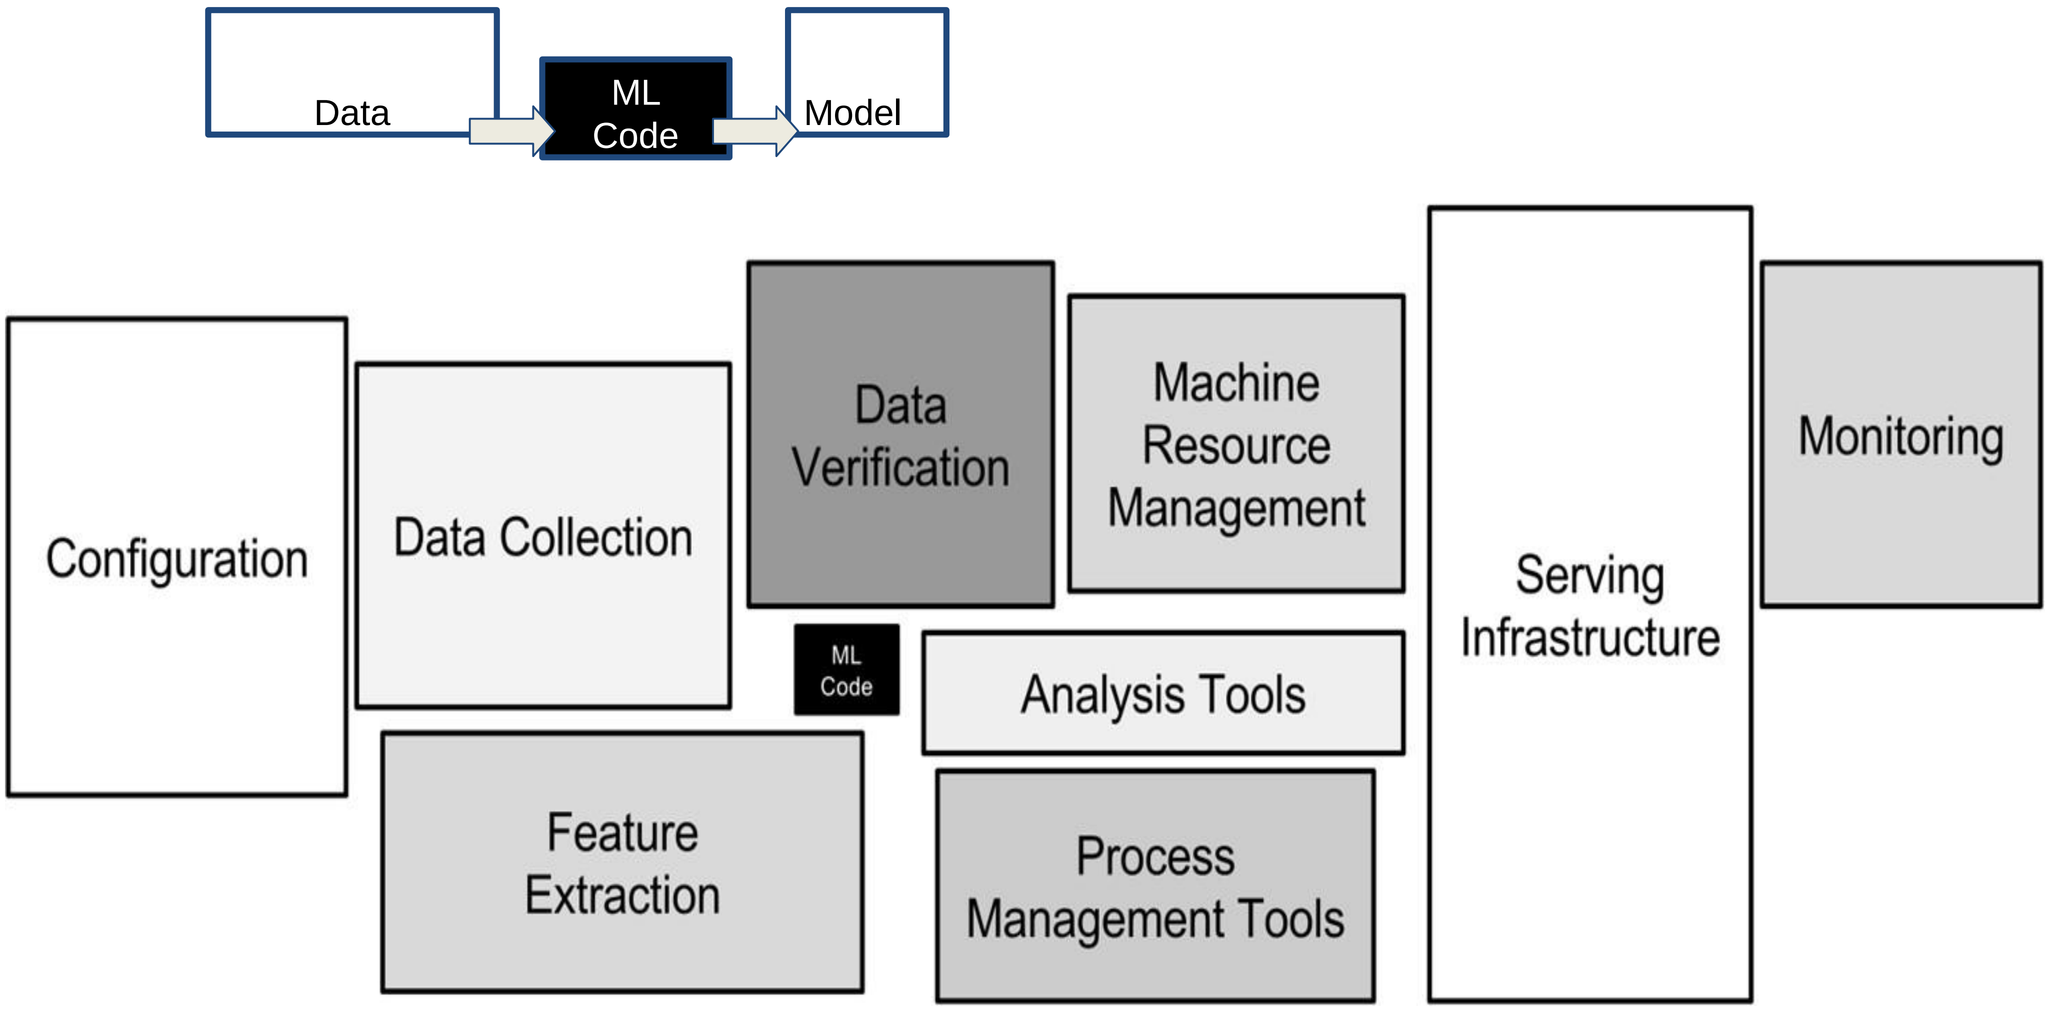

Hidden technical debt in Machine Learning Systems, Sculley et. al. (NIPS 2015)

What is Explainability and Explainable AI?

- An explanation is “a statement or account that makes something clear”
- Explainability is
  - “The ability to explain or to present in understandable terms to a human.” Finale Doshi-Velez and Been Kim in Towards A Rigorous Science of Interpretable Machine Learning (https://arxiv.org/abs/1702.08608)
  - “When you can stop asking why” Gilpin, et al in Explaining Explanations: An Approach to Evaluating Interpretability of Machine Learning (https://arxiv.org/abs/1806.00069)
- Explainable AI is an “AI systems that can explain their rationale to a human user, characterize their strengths and weaknesses, and convey an understanding of how they will behave in the future” DARPA XAI Program (https://www.darpa.mil/program/explainable-artificial-intelligence)

Explainable AI – The DARPA View

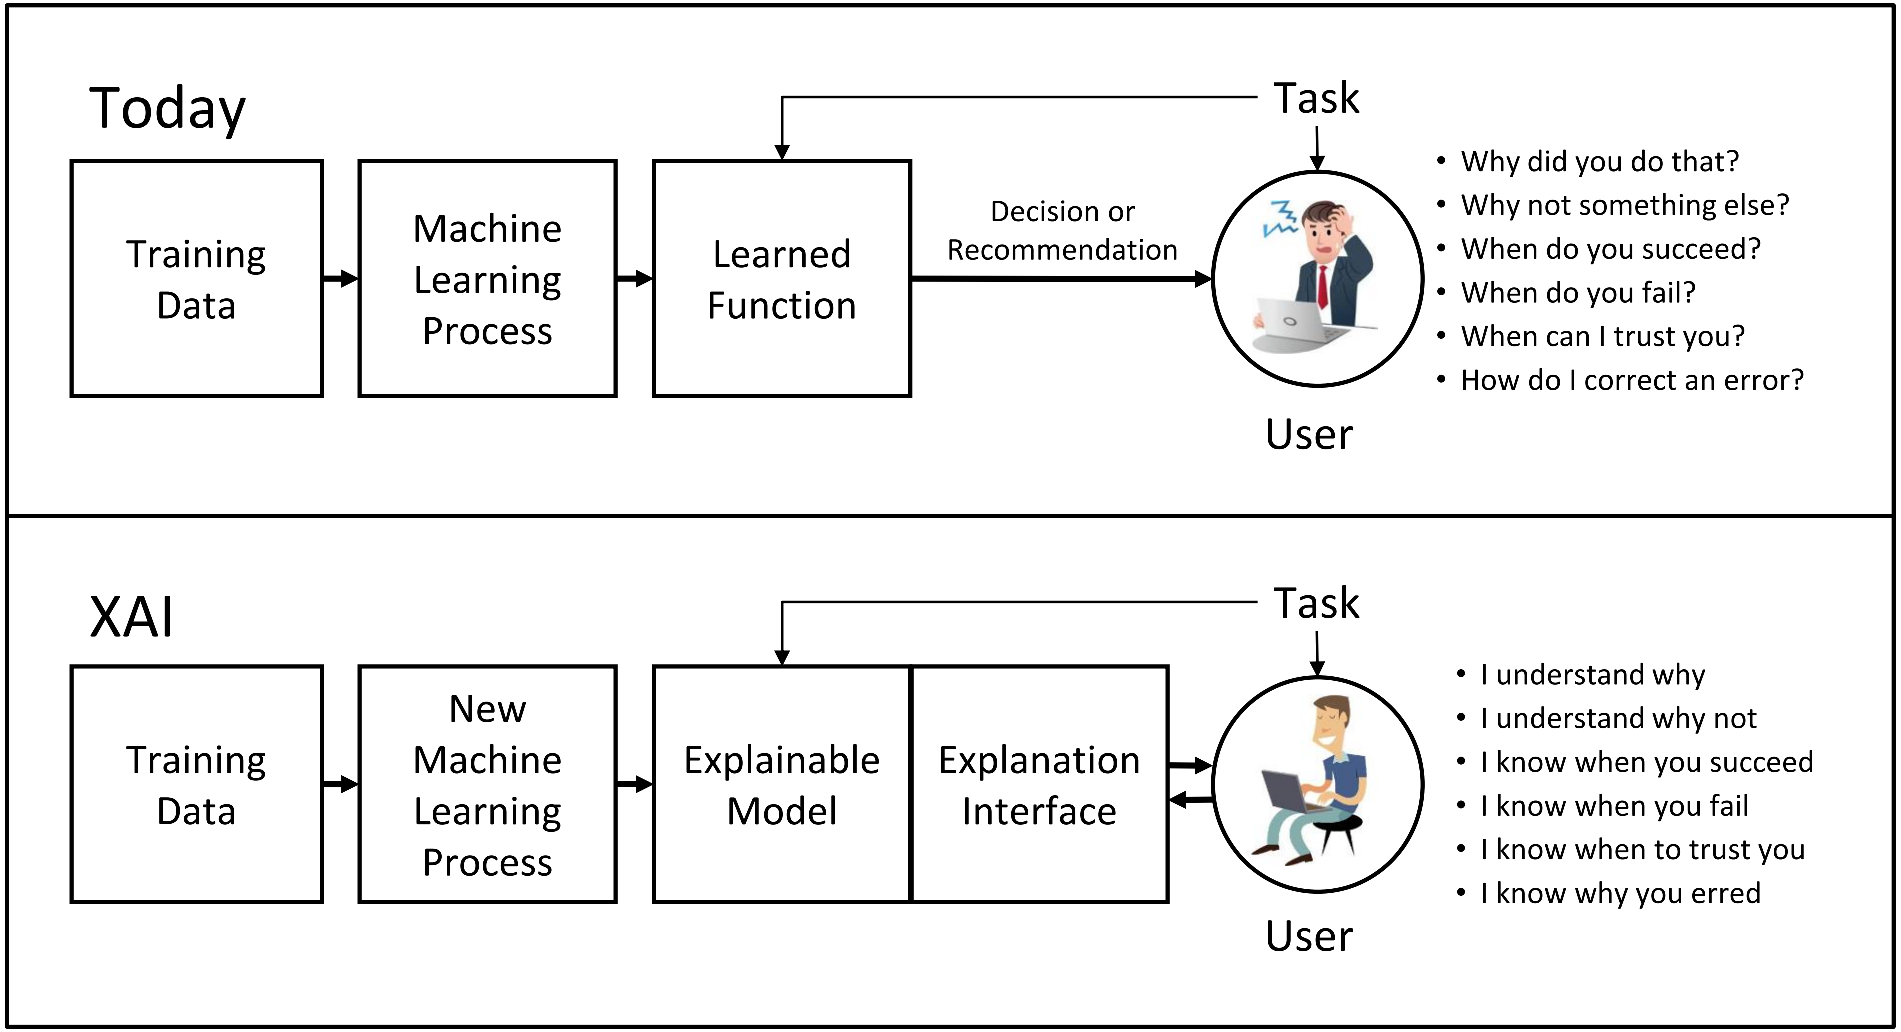

https://www.darpa.mil/attachments/XAIProgramUpdate.pdf

New Approach

Create a suite of machine learning techniques that produce more explainable models, while maintaining a high level of learning performance

Learning Techniques (today)

Explainability (notional)

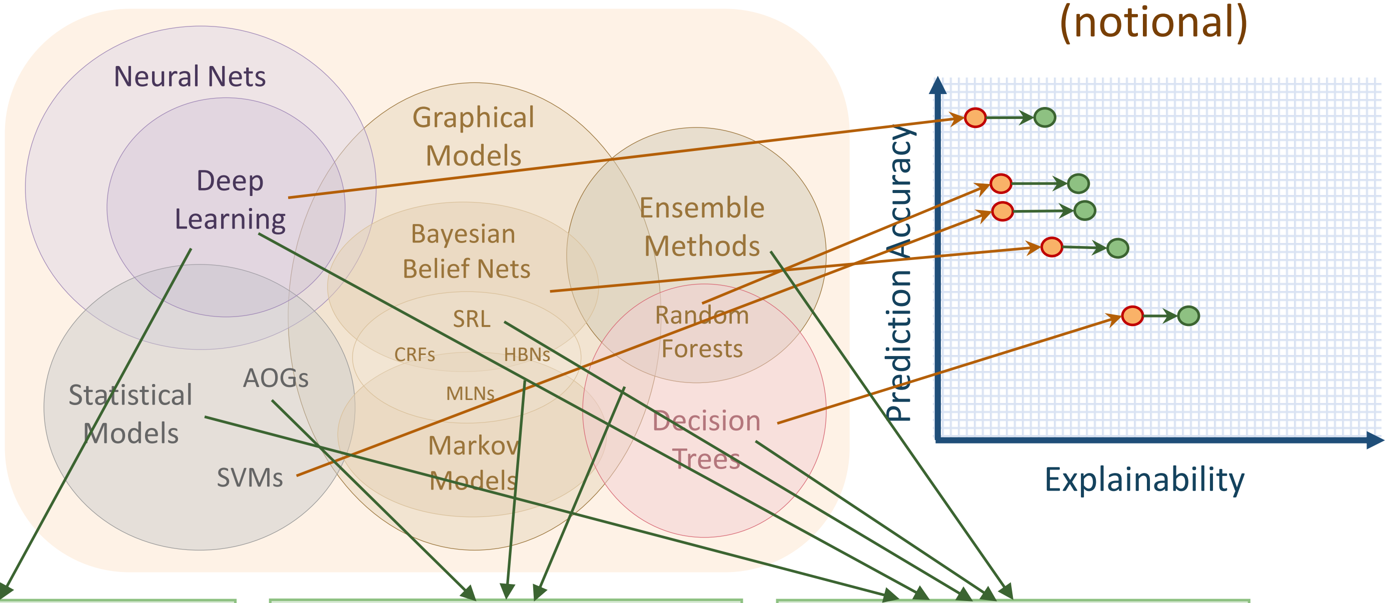

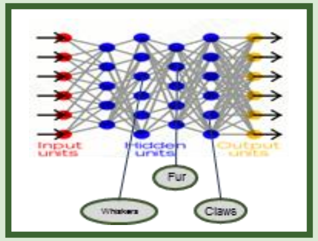

Deep Explanation

Modified deep learning techniques to learn explainable features

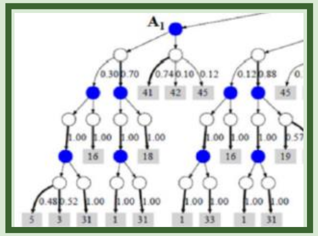

Interpretable Models

Techniques to learn more structured, interpretable, causal models

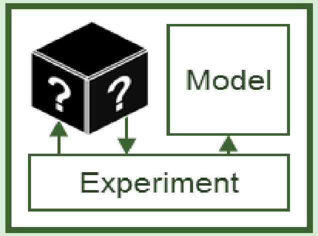

Model Induction

Techniques to infer an explainable model from any model as a black box

LIME (Local Interpretable Model-agnostic Explanations)

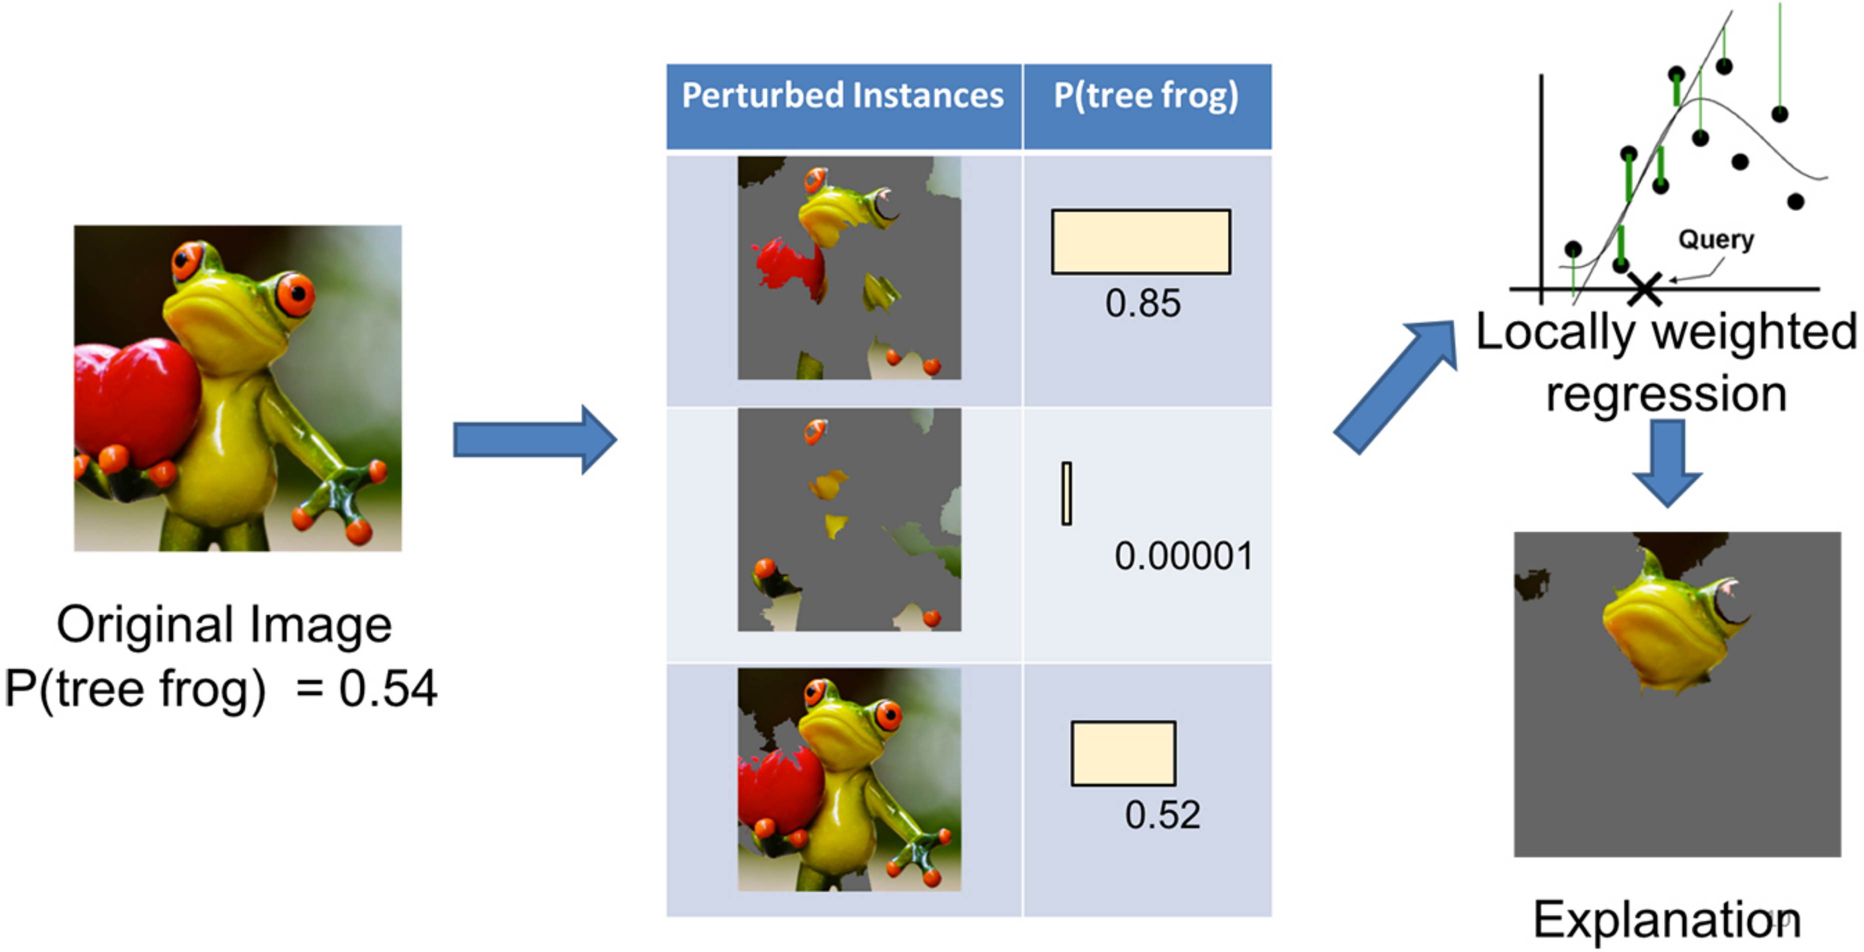

https://www.oreilly.com/learning/introduction-to-local-interpretable-model-agnostic-explanations-lime

Safe Autonomous Systems / AI

If things can go wrong they probably will!

This implies the need for continual monitoring of an autonomous system and its environment in a principled, contextual, task specific manner which can be specified by the system itself!

<u>Reasoning over Uncertainty</u>

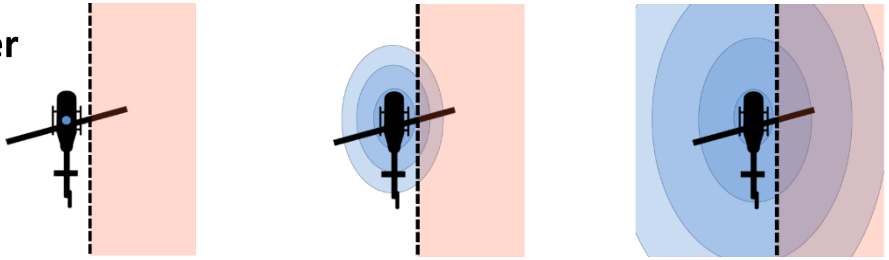

collision: false   Pr(collision) = 0.1   Pr(collision) = 0.4

<u>Reasoning over Predictions</u>

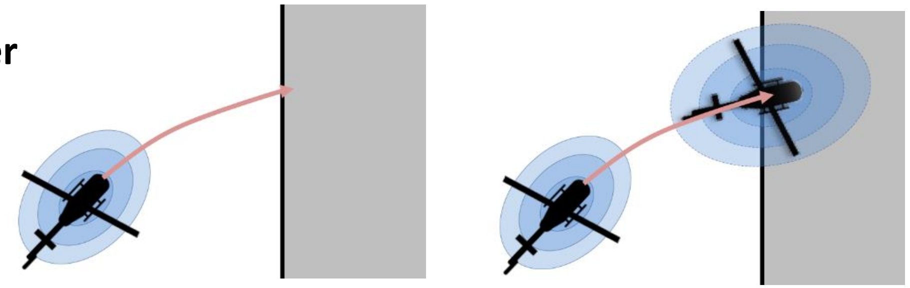

Pr(collision now) = 0.0...   Pr(collision soon) = 0.5

Probabilistic Predictive Stream Reasoning [Tiger and Heintz TIME 2016, IJAR review]

Reinforcement Learning

Reinforcement Learning Basic Concept

- *Reinforcement Learning is learning what to do – how to map situations to actions – so as to maximum a numerical reward.*

  Reinforcement Learning: An introduction
  Sutton & Barto

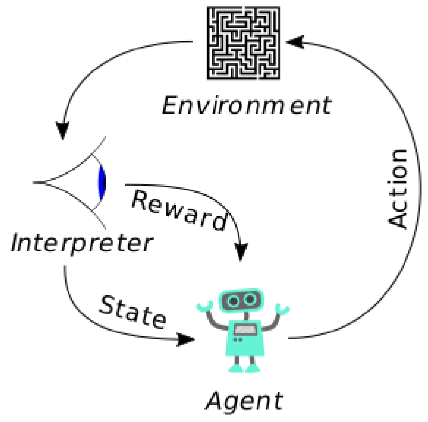

- Rather than learning from explicit training data, or discovering patterns in static data, reinforcement learning discovers the best option (highest reward) from trial and error.
- Inverse Reinforcement Learning
  - Learn reward function by observing an expert
  - “Apprenticeship learningapprenticeship learning“
  - E.g. Abbeel et al. *Autonomous Helicopter Aerobatics through Apprenticeship Learning*

A Reinforcement Learning Problem

- The environment
- The reinforcement function *r(s,a)*
  - Pure delay reward and avoidance problems
  - Minimum time to goal
  - Games
- The value function *V(s)*
  - Policy $\pi: S \rightarrow A$
  - Value $V^{\pi}(s) := \Sigma_i \gamma^i r_{t+i}$
- Find the optimal policy π^* that maximizes $V^{\pi^*}(s)$ for all states *s*.

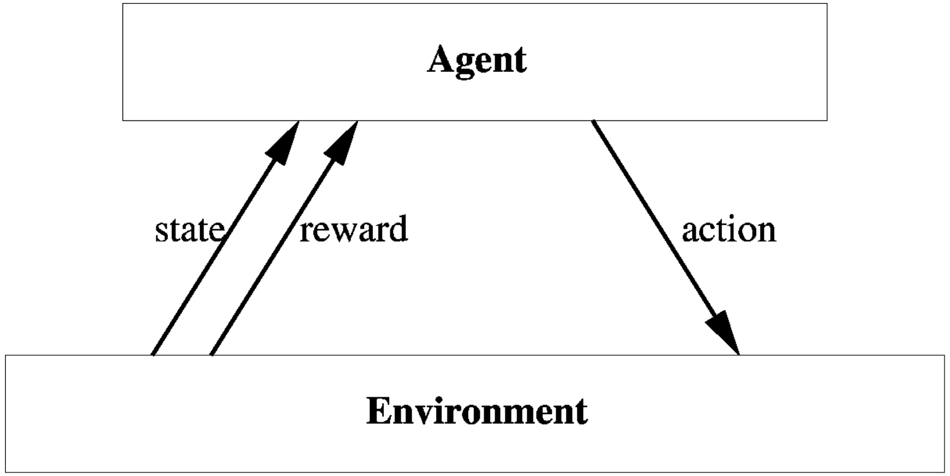

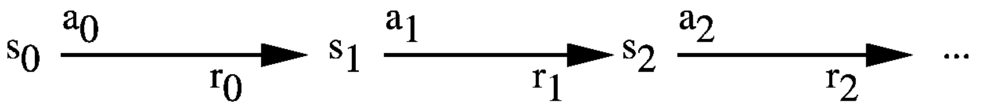

Goal: Learn to choose actions that maximize

$r_0 + \gamma r_1 + \gamma^2 r_2 + \ldots$, where $0<\gamma<1$

RL Value Function - Example

A minimum time to goal world

0	-14	-20	-22
-14	-18	-22	-20
-20	-22	-18	-14
-22	-20	-14	0

Value function for random movement

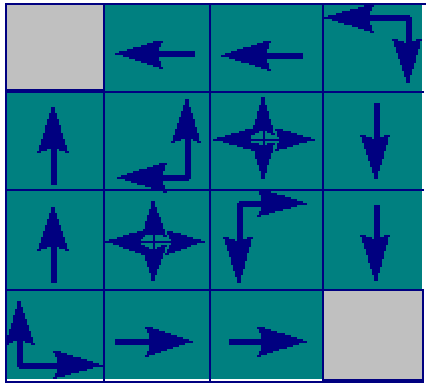

Optimal policy

0	-1	-2	-3
-1	-2	-3	-2
-2	-3	-2	-1
-3	-2	-1	0

Optimal value function

Markov Decision Processes

Assume:

- finite set of states S, finite set of actions A
- at each discrete time agent observes state $s_t \in S$ and chooses action $a_t \in A$
- then receives immediate reward r_t
- and state changes to s_{t+1}
- Markov assumption: $s_{t+1} = \delta(s_t, a_t)$ and $r_t = \mathrm{r}(s_t, a_t)$
  - i.e. r_t and s_{t+1} depend only on current state and action
  - functions δ and r may be non-deterministic
  - functions δ and r not necessarily known to the agent

MDP Example

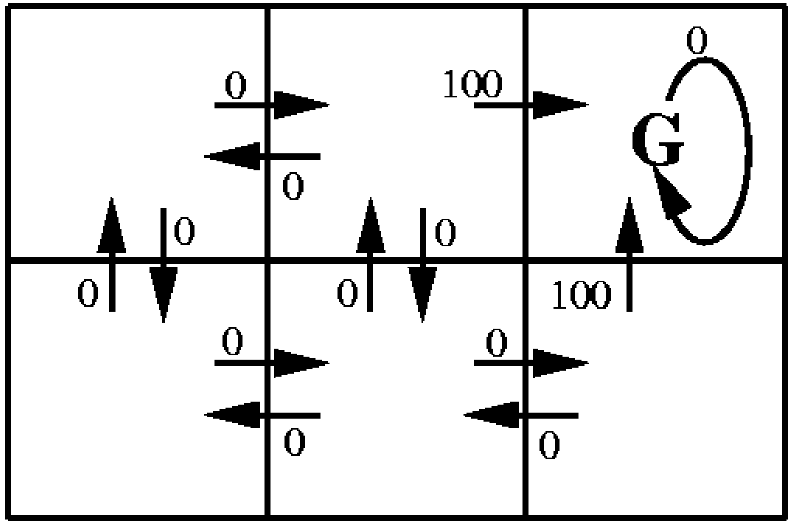

r(s,a)

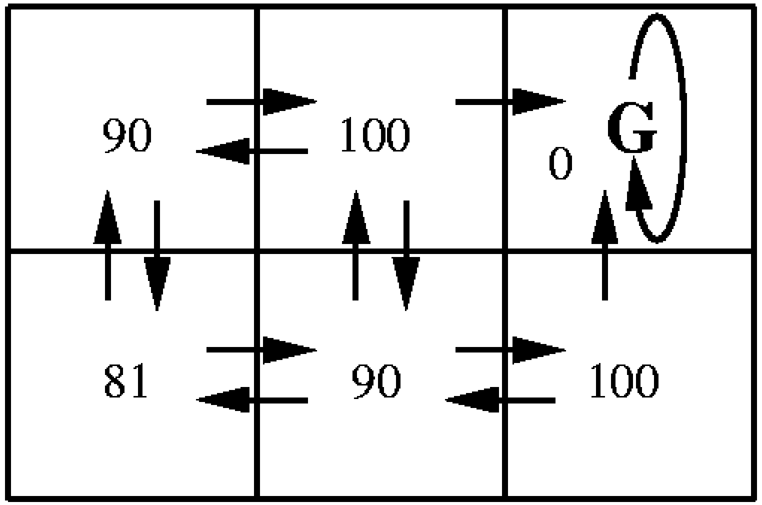

V(s)*

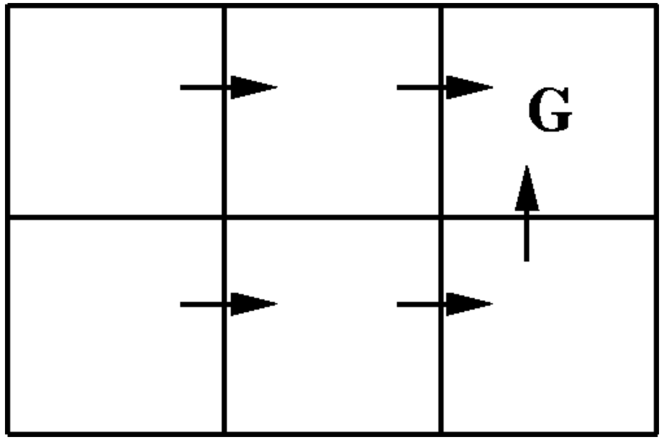

An optimal policy

The Q-Function

Optimal policy:

- $\pi^*(s) = \text{argmax}_a[r(s,a) + \gamma V^*(\delta(s,a))]$
- Doesn't work if we don't know r and δ.

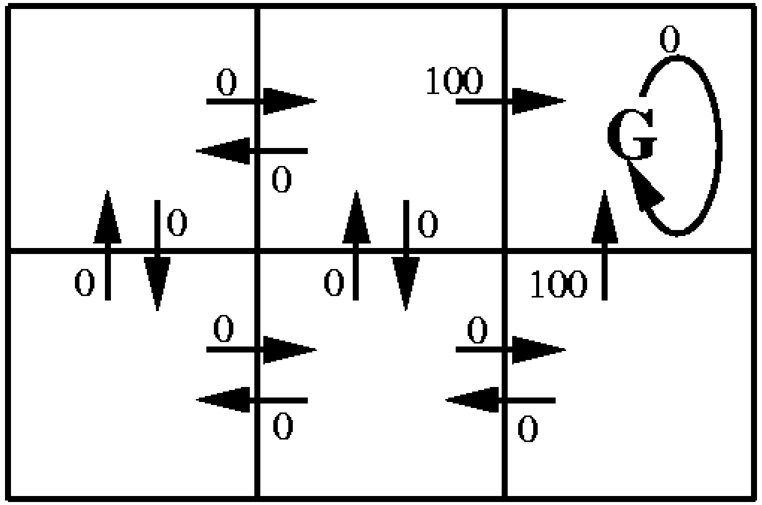

r(s,a)

The Q-function:

- $Q(s,a) := r(s,a) + \gamma V^*(\delta(s,a))$
- $\pi^*(s) = \text{argmax}_a Q(s,a)$

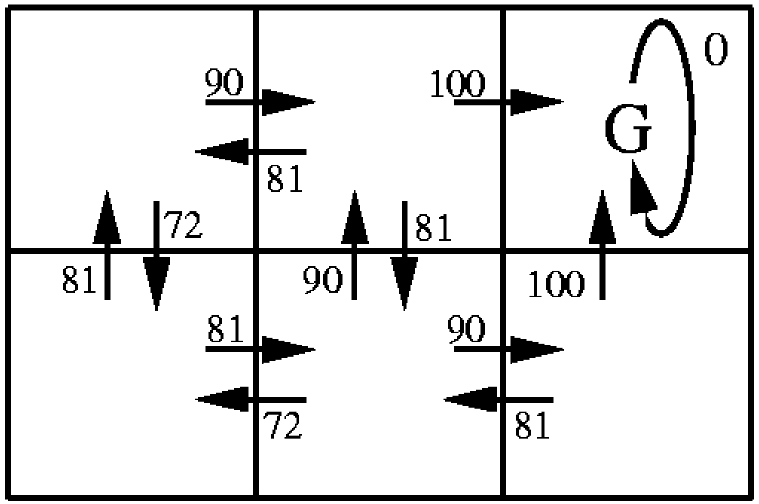

Q(s,a)

The Q-Function

- Note Q and V* closely related:
  $V^*(s) = \max_{a'} Q(s,a')$
- Therefore Q can be written as:

  $$Q(s_t, a_t) := r(s_t, a_t) + \gamma V^*(\delta(s_t, a_t)) =$$
  $$r(s_t, a_t) + \gamma \max_{a'} Q(s_{t+1}, a')$$

- If $Q^\wedge$ denote the current approximation of Q then it can be updated by:

  $$Q^\wedge(s,a) := r + \gamma \max_{a'} Q^\wedge(s',a')$$

Q-Learning for Deterministic Worlds

For each s, a initialize table entry $\hat{Q}(s,a) := 0$.

Observe current state s.

Do forever:

1. Select an action a and execute it
2. Receive immediate reward r
3. Observe the new state s'
4. Update the table entry for $\hat{Q}(s,a)$:
   $$\hat{Q}(s,a) := r + \gamma \max_{a'} \hat{Q}(s',a')$$
5. $s := s'$

Q-Learning Example

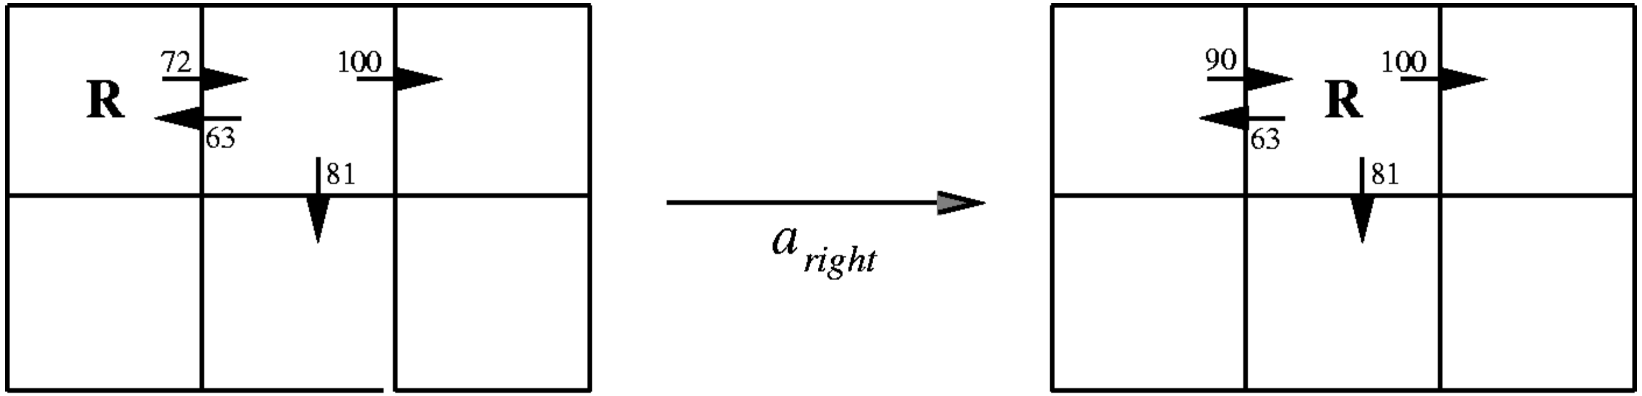

Initial state: s_1 Next state: s_2

$$\hat{Q}(s_1, a_{right}) := r + \gamma \max_{a'} \hat{Q}(s_2, a')$$
$$:= 0 + 0.9 \max\{63, 81, 100\}$$
$$:= 90$$

Q-Learning Continued

- Exploration
  - Selecting the best action
  - Probabilistic choice
- Improving convergence
  - Update sequences
  - Remember old state-action transitions and their immediate reward
- Non-deterministic MDPs
- Temporal Difference Learning

Reinforcement Learning – Neural Networks as Function Approximators

- To tackle a high-dimensional state space or continous states we can use a neural network as function approximator
- Lunar Lander experiment
  - 8 continous/discrete states
    - XY-Pos, XY-Vel, Rot, Rot-rate, Leg1/Leg2 ground contact
  - 4 discrete actions
    - Left thrust
    - Right thrust
    - Main engine thrust
    - NOP
  - Rewards
    - Move from top to bottom of the screen (+ ~100-140)
    - Land between the posts (+100)
    - Put legs on ground (+10 per leg)
  - Penalties
    - Using main engine thrust (-0.3 per frame)
    - Crashing (-100)
- Solved using Stochastic Policy Gradients

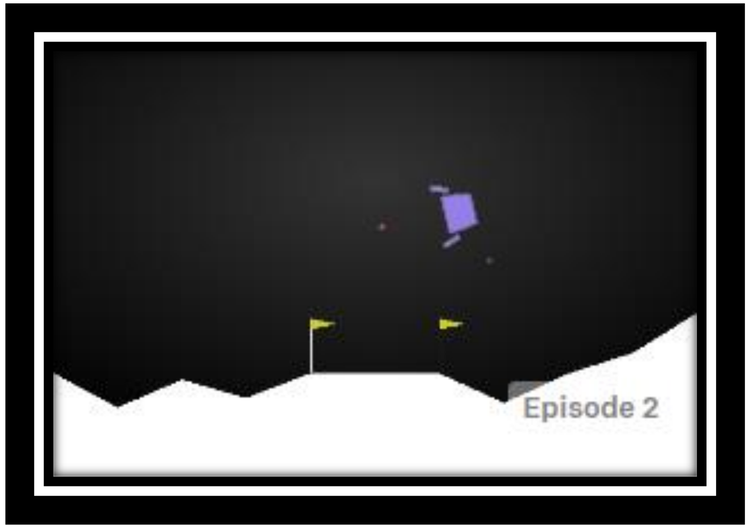

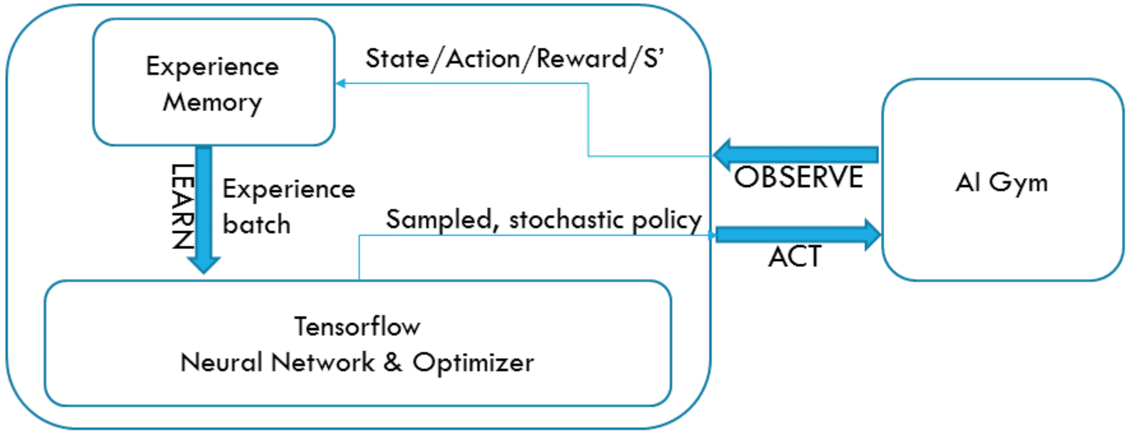

Reinforcement Learning Neural Networks as Function Approximators

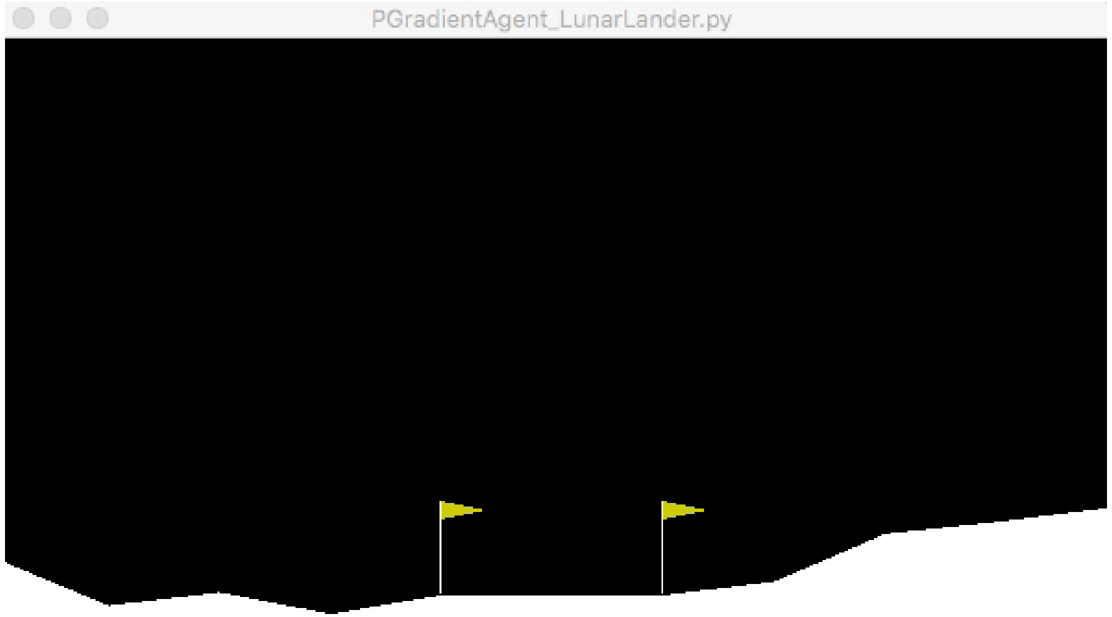

Reinforcement Learning Basic Concepts

- Value-Based:
  - Learn value function
  - Implicit policy (e.g. greedy selection)
  - Example: Deep Q Networks (DQN)
- Policy-Based:
  - No value function
  - Learn explicit (stochastic) policy
  - Example: Stochastic Policy Gradients
- Actor-Critic:
  - Learn value function
  - Learn policy using value function
  - Example: Asynchronous Advantage Actor Critic (A3C)

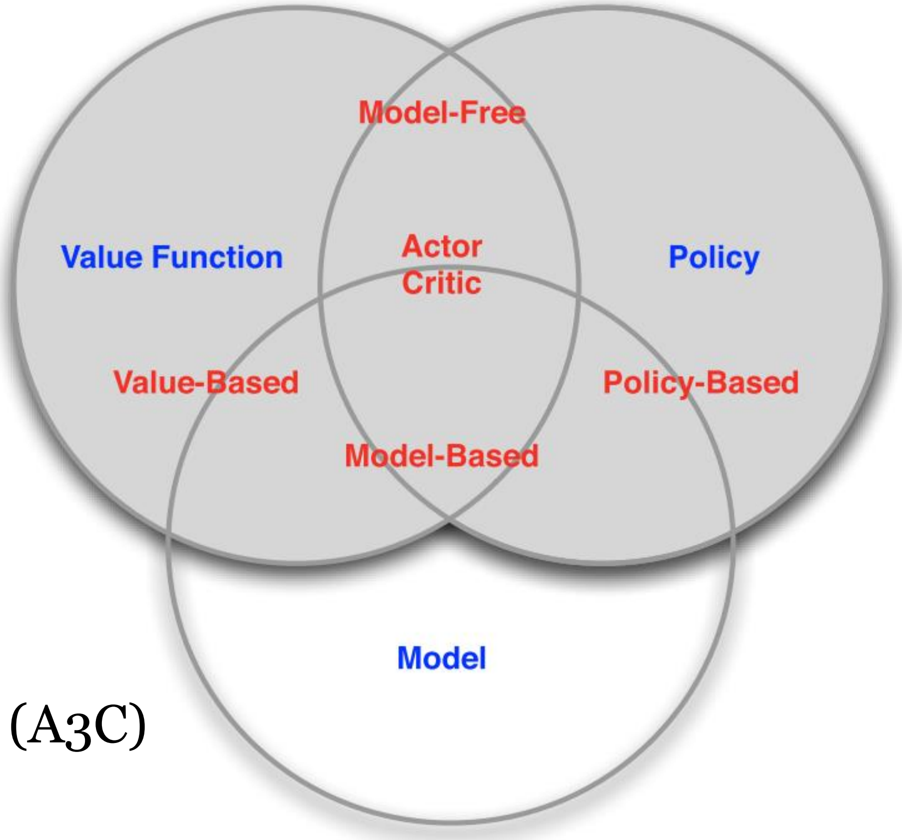

OpenAI Gym Functionality

- Algorithms
  - Imitate computations
- Atari
  - Reach high scores in Atari 2600 games
- Box2D
  - Continuous/Discrete control tasks in Box2D simulator
- Classic Control
  - Control theory problems from classic RL litterature
- MuJoCo
  - Continuous control tasks
- Robotics
  - Simulated goal-based tasks for fetch and shadow hand robots
- Toy text
  - Simple text environments

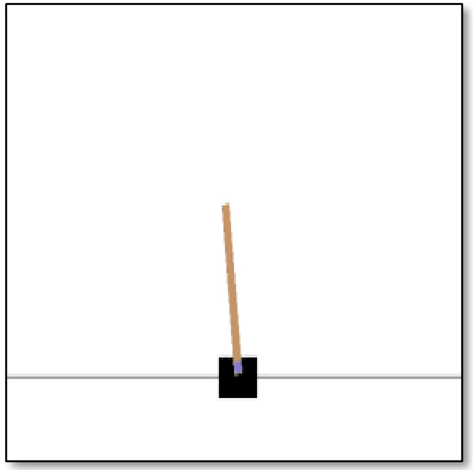

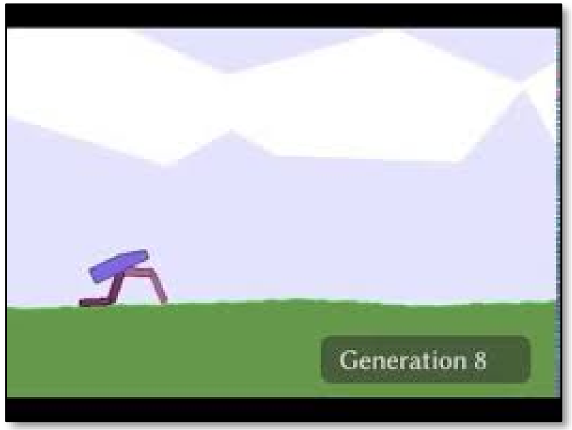

LVC Simulation for Improved Training Efficiency

- This project will use agent-based simulation to address the growing need for efficient and effective pilot training solutions for fighter aircraft, in an LVC context:
  - Lower costs
  - Improve availability
  - Realize more complex scenarios to improve training value
- We will develop machine learning techniques that allow agents to learn in complex environments:
  - Mixed cooperative and competitive multi-agent scenarios
  - Multiple conflicting objectives
  - Partial observability
  - Sparse rewards

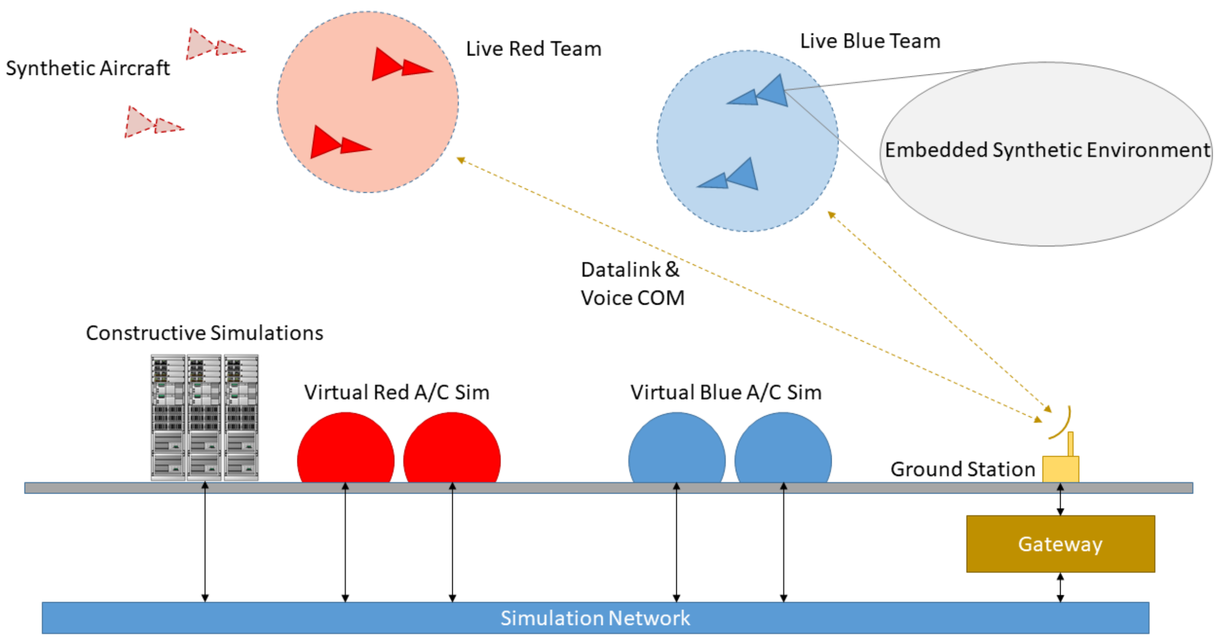

LINKÖPINGS UNIVERSITET

LVC Simulation for Improved Training Efficiency

- Similar challenges as for real-time strategy games (e.g. Dota and StarCraft):
  - Long time horizons for decision making:
    - Dota ~20000 moves per game (45 min), Go ~150 moves per game, chess ~40 moves per game
  - Partially observed state, e.g. sensor limitations and EW
  - Complex observation and action spaces
- As a result, exploration takes longer time:
  - E.g. training of OpenAI Five: 128,000 CPU cores for simulation rollouts and 256 GPUs for training of neural network model (~180 years of game play experience per day)

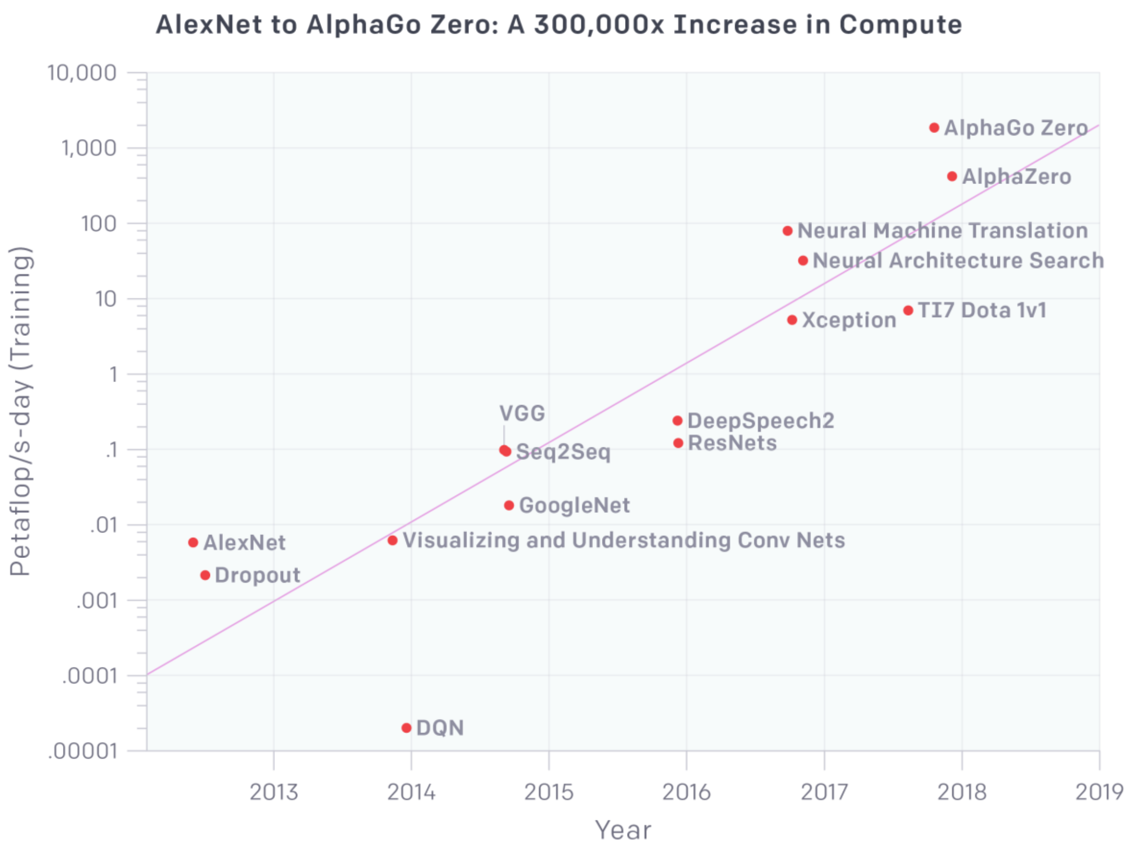

Multi-Objective Reinforcement Learning (MORL)

- Many real-world tasks may present an agent with multiple, possibly conflicting objectives:
  - Time
  - Safety
  - Resource consumption
- Multi-Objective Reinforcement Learning allows an agent to learn how to prioritize among objectives at runtime
- Possible to create diverse populations of agents, or adapt agents to time-varying user needs, e.g. difficulty level or training session contents
- Training goals can also be considered by agents

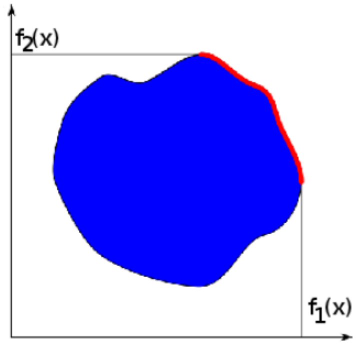

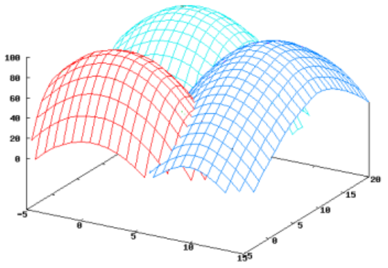

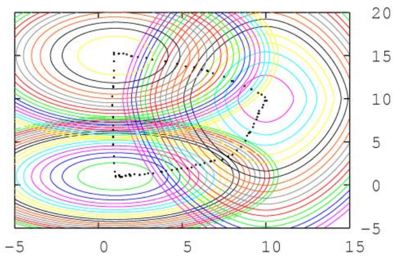

LINKÖPINGS UNIVERSITET

J. Källström and F.Heintz, Tunable Dynamics in Agent-Based Simulation using Multi-Objective Reinforcement Learning, AAMAS Adaptive and Learning Agents Workshop 2019.

Conclusions

Algorithmic, Knowledge-Based and Learning-Based AI

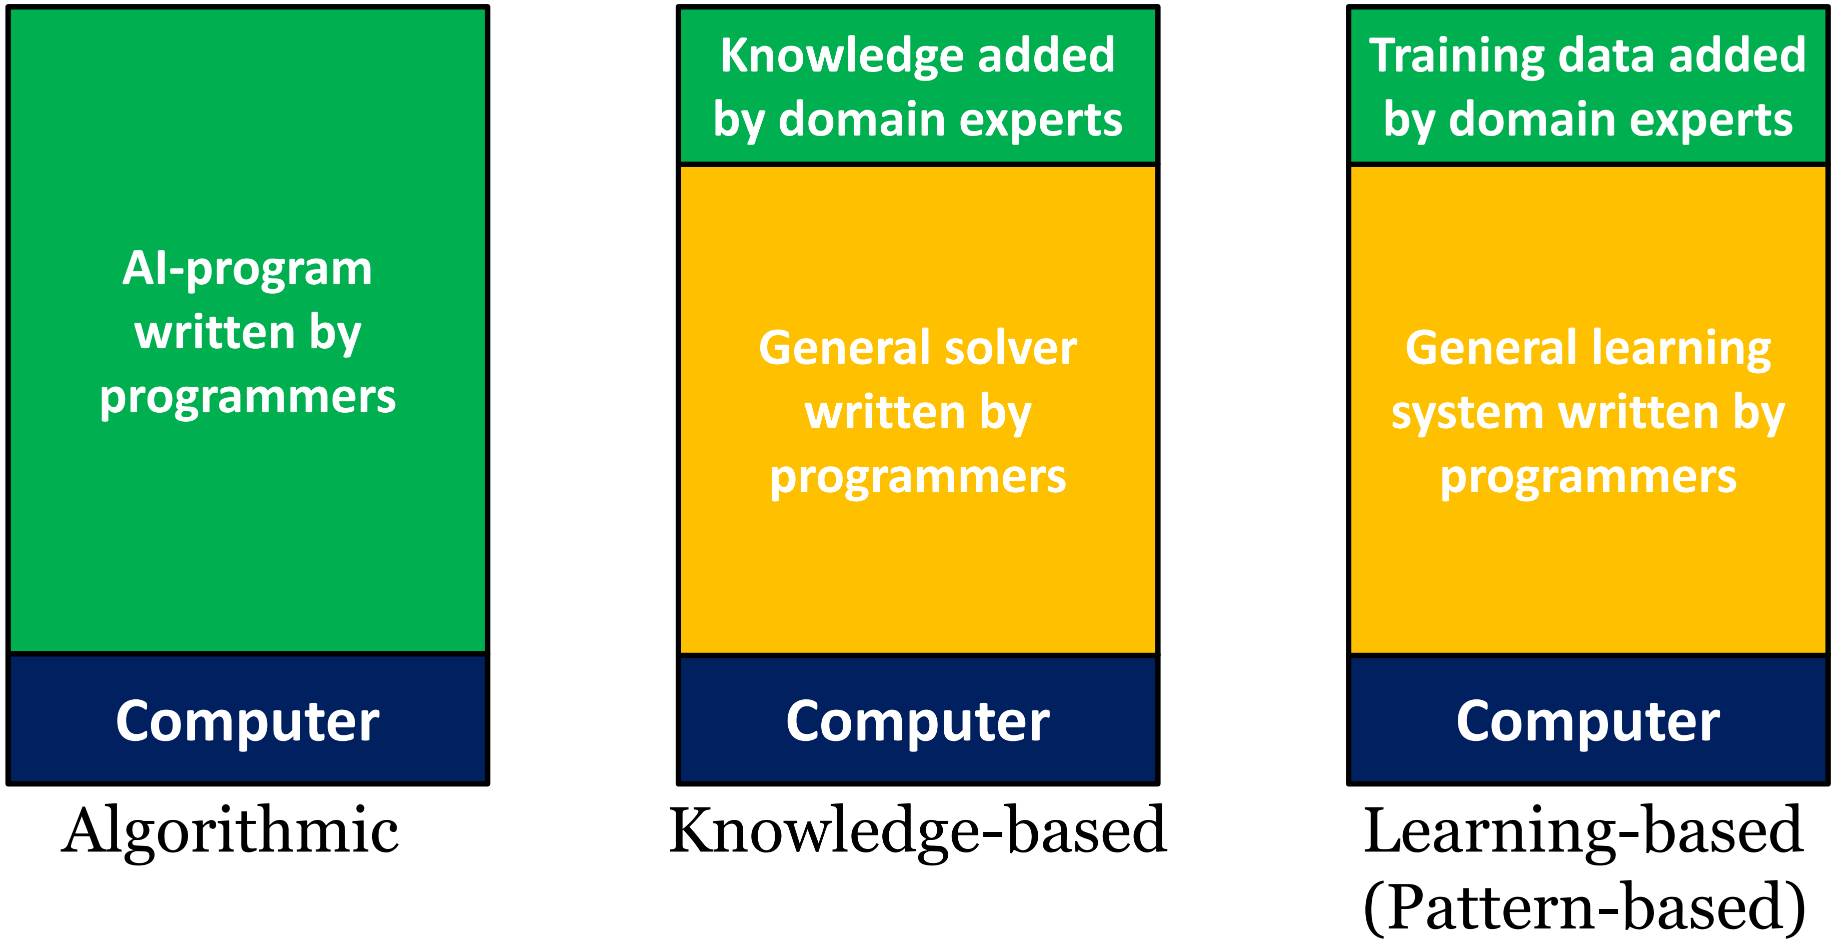

Algorithms vs Machine Learning

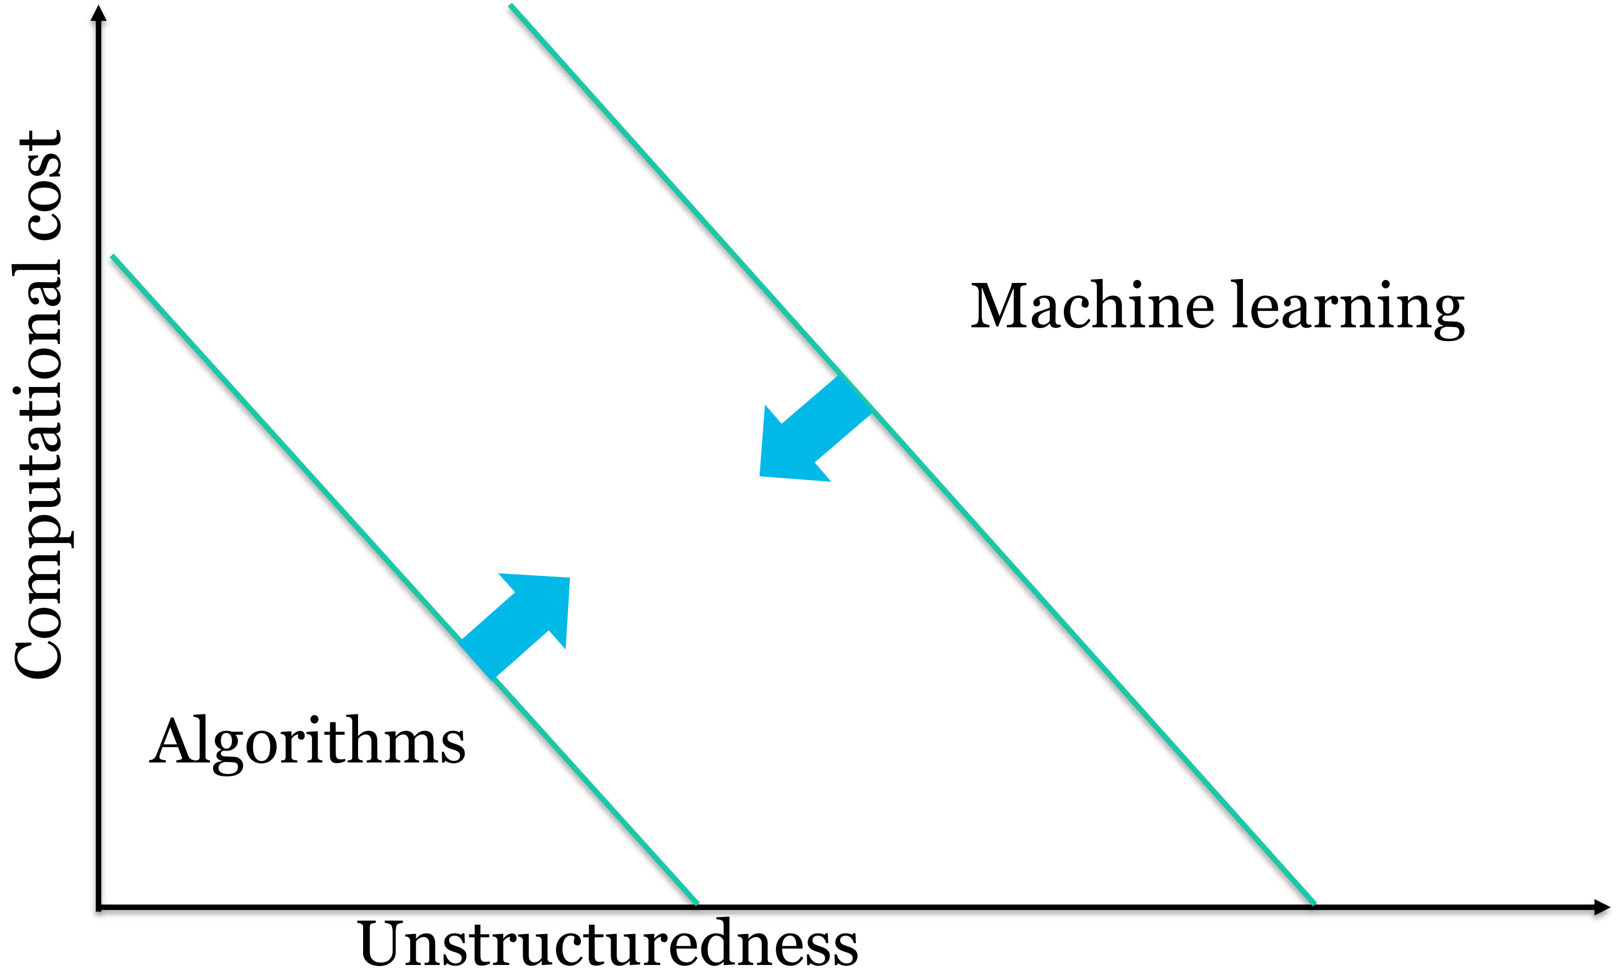

Human and Computational Thinking

Figure 1: A Comparison of System 1 and System 2 Thinking

System 1 "Fast"	**System 2** "Slow"
DEFINING CHARACTERISTICS	DEFINING CHARACTERISTICS
Unconscious Effortless Automatic	Deliberate and conscious Effortful Controlled mental process
WITHOUT self-awareness or control	WITH self-awareness or control
"What you see is all there is."	Logical and skeptical
ROLE	ROLE
Assesses the situation Delivers updates	Seeks new/missing information Makes decisions

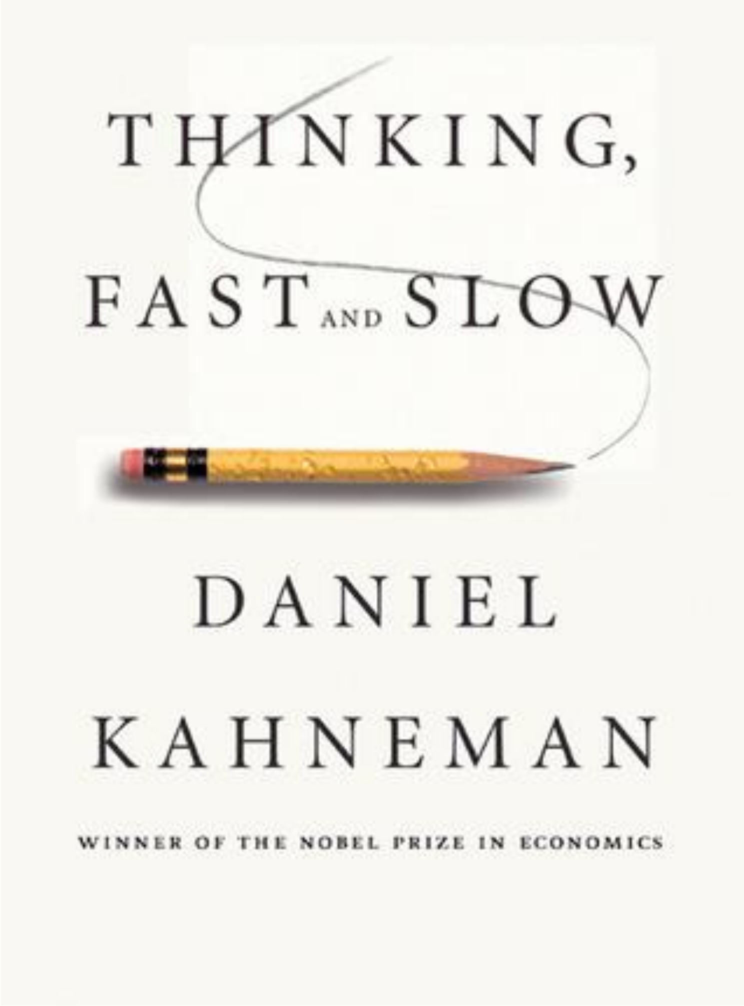

Pure Logic **Pure Learning**

- Slow thinking: deliberative, cognitive, model-based, extrapolation
- Amazing achievements until this day
- "*Pure logic is brittle*"
  noise, uncertainty, incomplete knowledge, …

https://web.cs.ucla.edu/~guyvdb/slides/ComputersAndThought.pdf

Pure Logic **Pure Learning**

- Fast thinking: instinctive, perceptive, model-free, interpolation
- Amazing achievements recently
- “*Pure learning is brittle*”

  bias, algorithmic fairness, interpretability, explainability, adversarial attacks, unknown unknowns, calibration, verification, missing features, missing labels, data efficiency, shift in distribution, general robustness and safety

  fails to incorporate a sensible model of the world

https://web.cs.ucla.edu/~guyvdb/slides/ComputersAndThought.pdf

The Way Forward

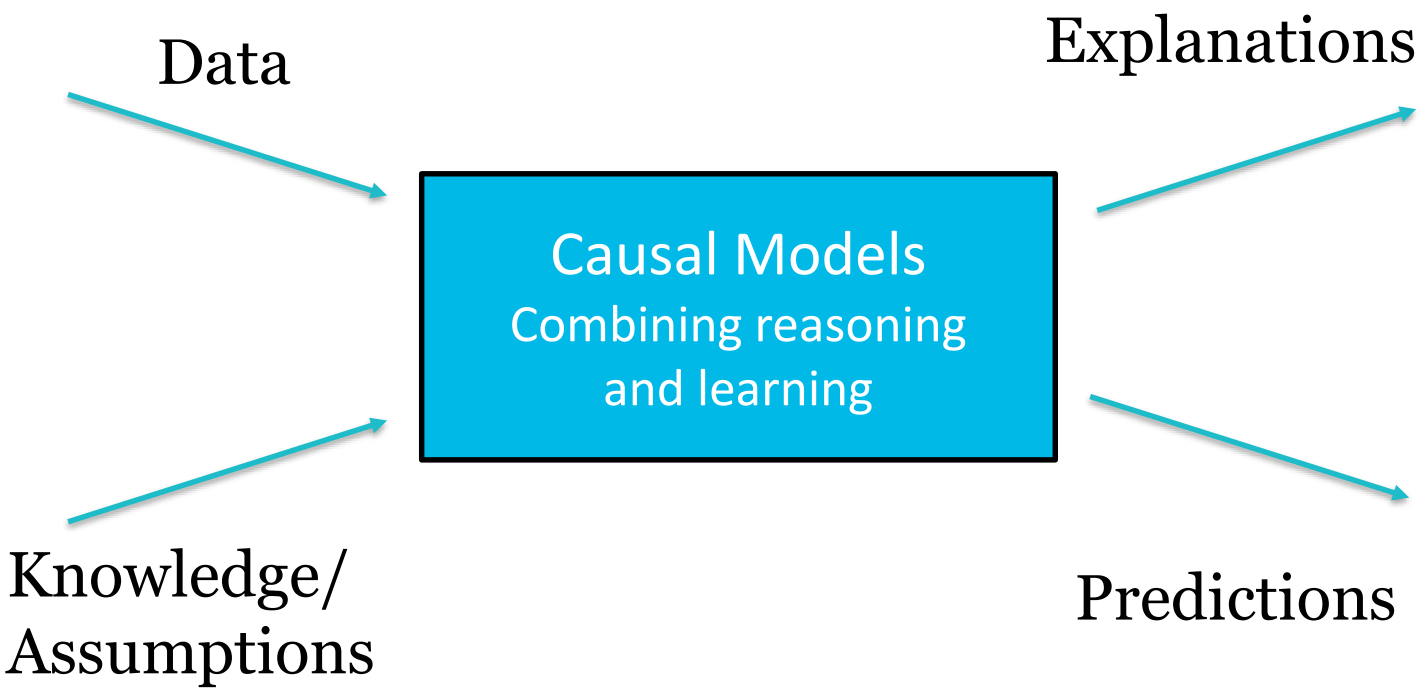

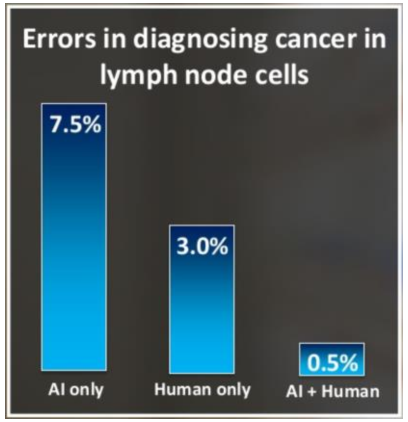

<

<

"Weak human + machine + superior process was greater than a strong computer and, remarkably, greater than a strong human + machine with inferior process."

Garry Kasparov

Basera medicinska beslut på all tillgänglig information? Beslutsstöd behövs!

LINKÖPINGS UNIVERSITET

Illustration gjord av Erik Sundvall, Region Östergötland & LiU. Delvis baserad på: Evidence-Based Medicine and the Changing Nature of Healthcare: Workshop Summary (IOM Roundtable on Evidence-Based Medicine) Mark B. McClellan, Michael McGinnis, Elizabeth G. Nabel, and LeighAnne M. Olsen, Institute of Medicine. ISBN: 0-309-11370-9 https://www.nap.edu/catalog/12041/evidence-based-medicine-and-the-changing-nature-of-health-care Fig 5-1. page 116

Computational Thinking

- Datalogiskt tänkande
- A **problem solving process** to describe, analyze, and solve problems such that computers can assist using techniques from computer science:
  - Give step-by-step instructions
  - Decompose problems into smaller parts
  - Find patterns
  - Create abstractions
  - Design algorithms

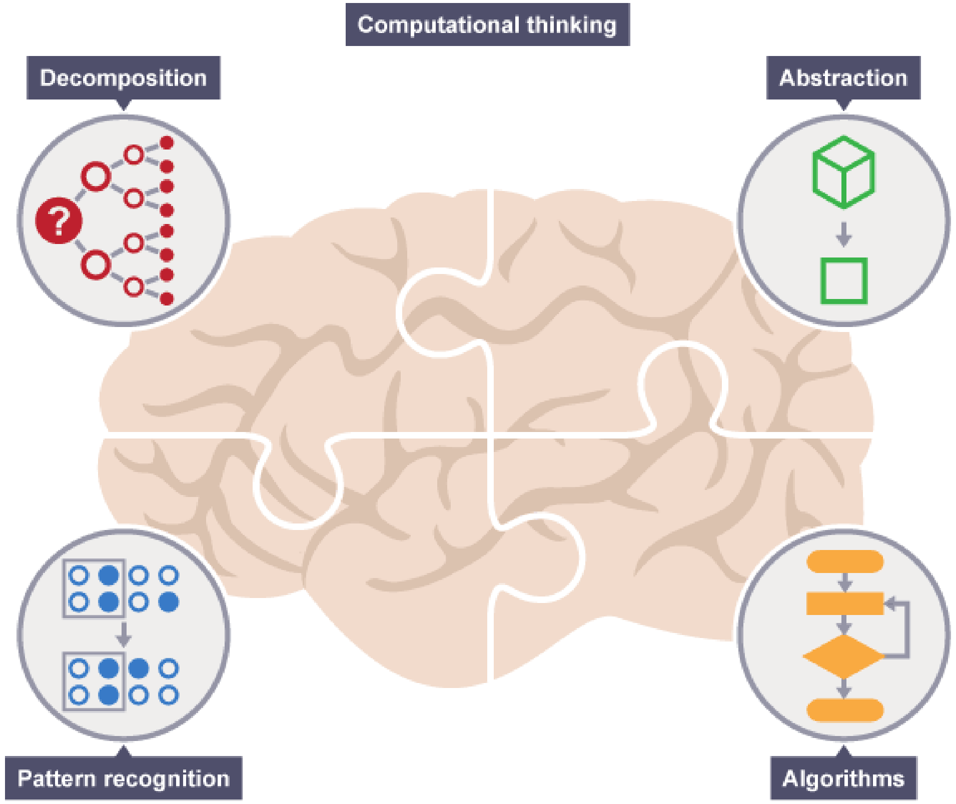

Ethics Guidelines for Trustworthy AI – Principles

4 Ethical Principles based on fundamental rights

Respect for human autonomy

Augment, complement and empower humans

Prevention of harm

Safe and secure. Protect physical and mental integrity.

Fairness

Equal and just distribution of benefits and costs.

Explicability

Transparent, open with capabilities and purposes, explanations

https://ec.europa.eu/digital-single-market/en/news/ethics-guidelines-trustworthy-ai

AI Innovation, Competence and Research in Sweden

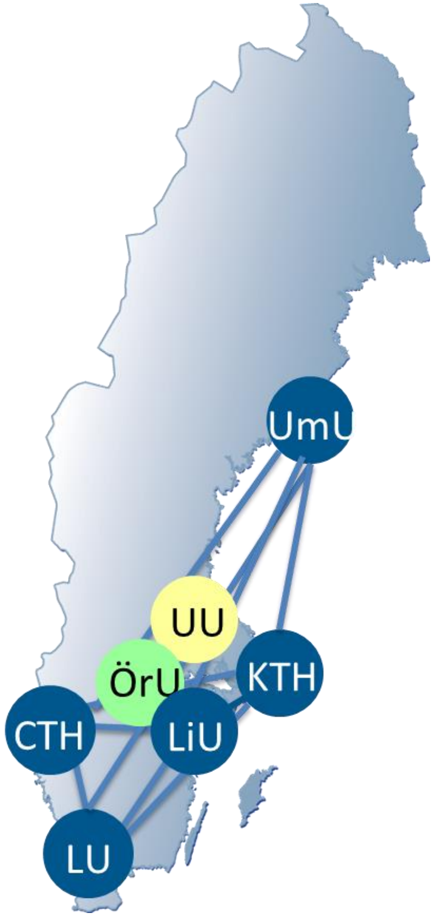

AI SUSTAINABILITY CENT

AI INNOVATION of Sweden

AI Competence of Swede

WASP | WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

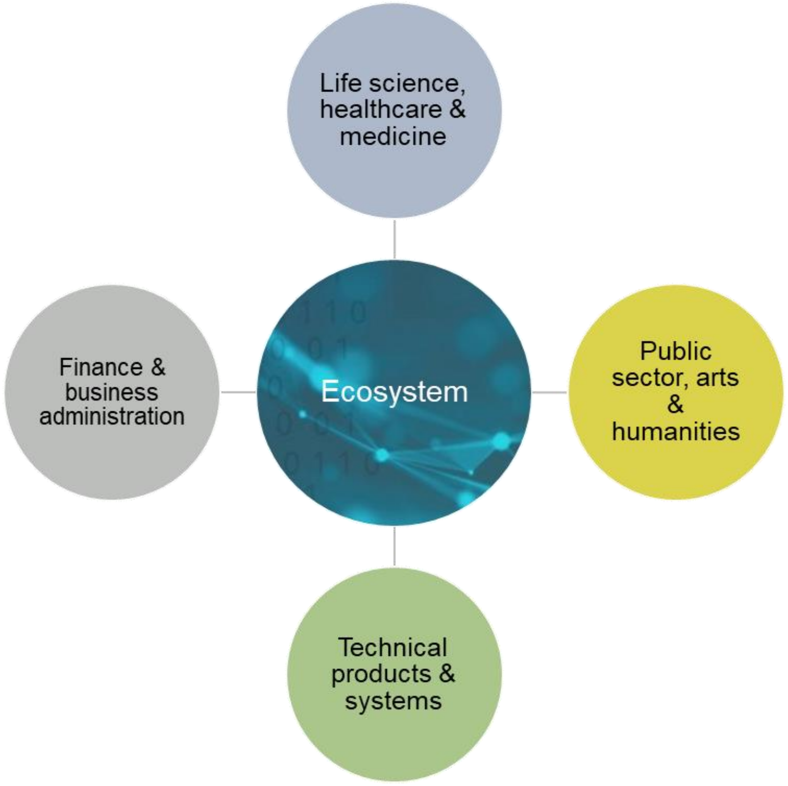

CHALMERS UNIVERSITY OF TECHNOLOGY · KTH VETENSKAP OCH KONST · LINKÖPINGS UNIVERSITET · LUNDS UNIVERSITET · UMEÅ UNIVERSITET · UPPSALA UNIVERSITET · ÖREBRO UNIVERSITY

Elements of AI

LINKÖPINGS UNIVERSITET

SWEDISH ARTIFICIAL INTELLIGENCE SOCIETY

Why is Artificial Intelligence Different

- Scale
- Speed
- Single-mindedness
- Optimization-based
- Cannot break the rules
- No needs
- No real consequences or “skin in the game”

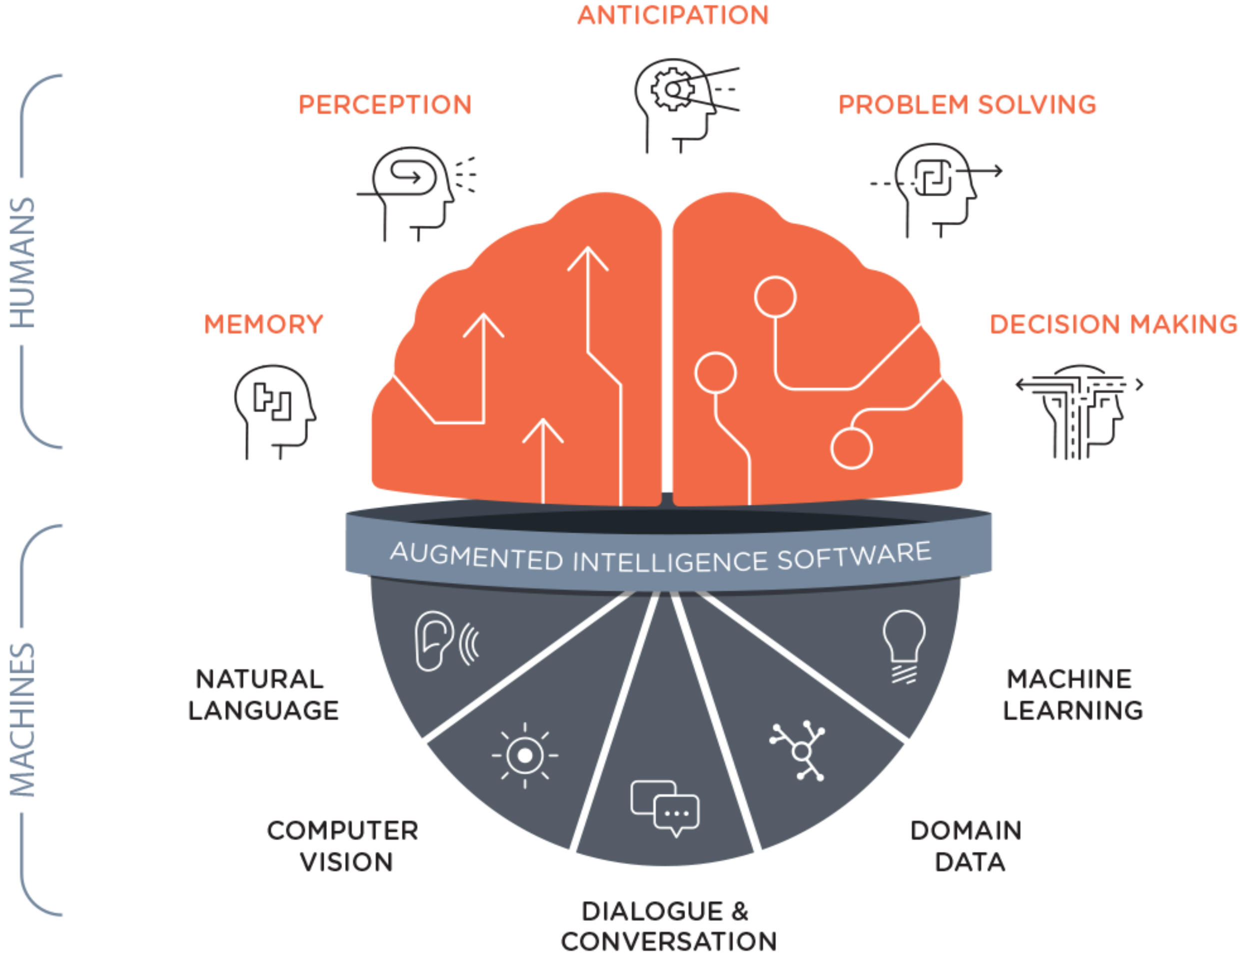

Prediction is hard, especially about the future

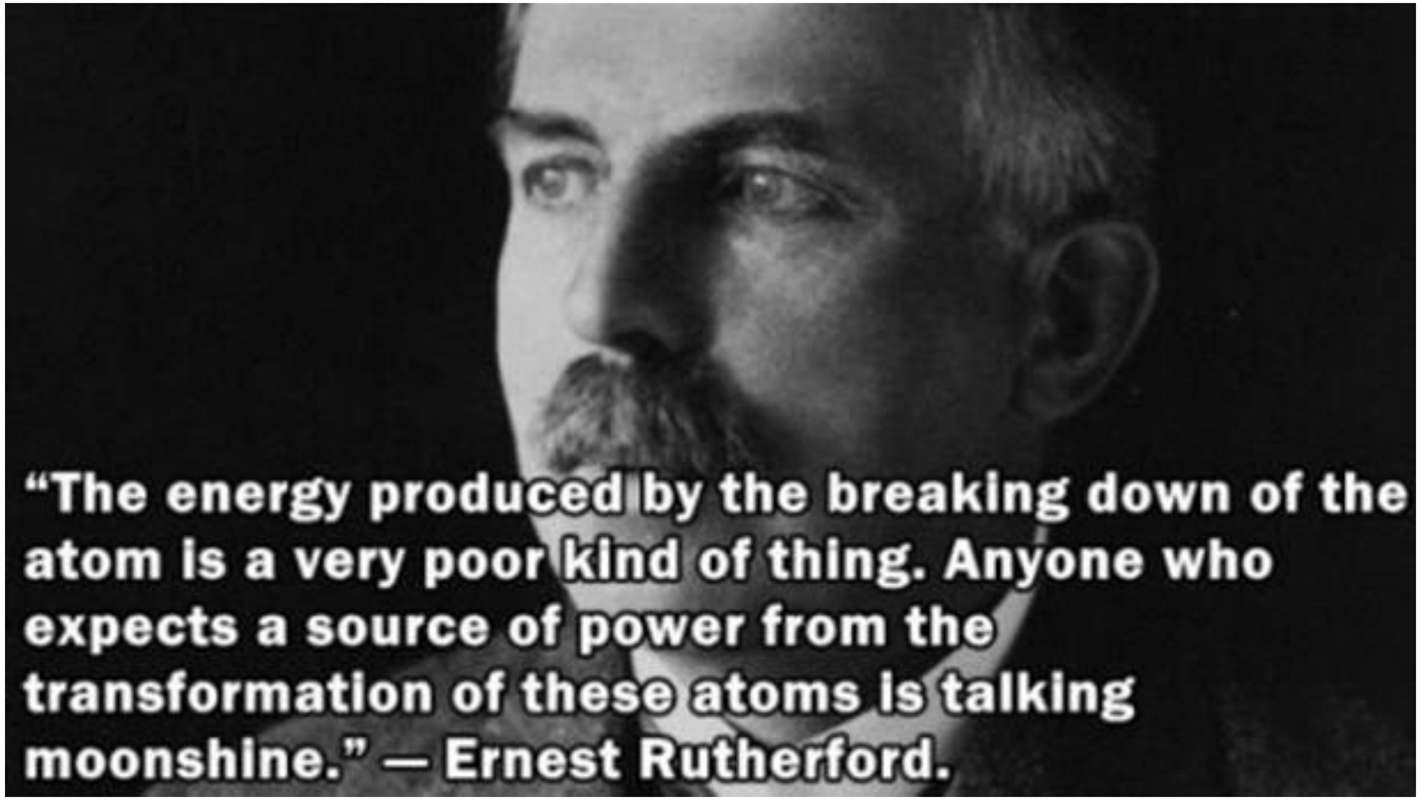

September 11, 1933

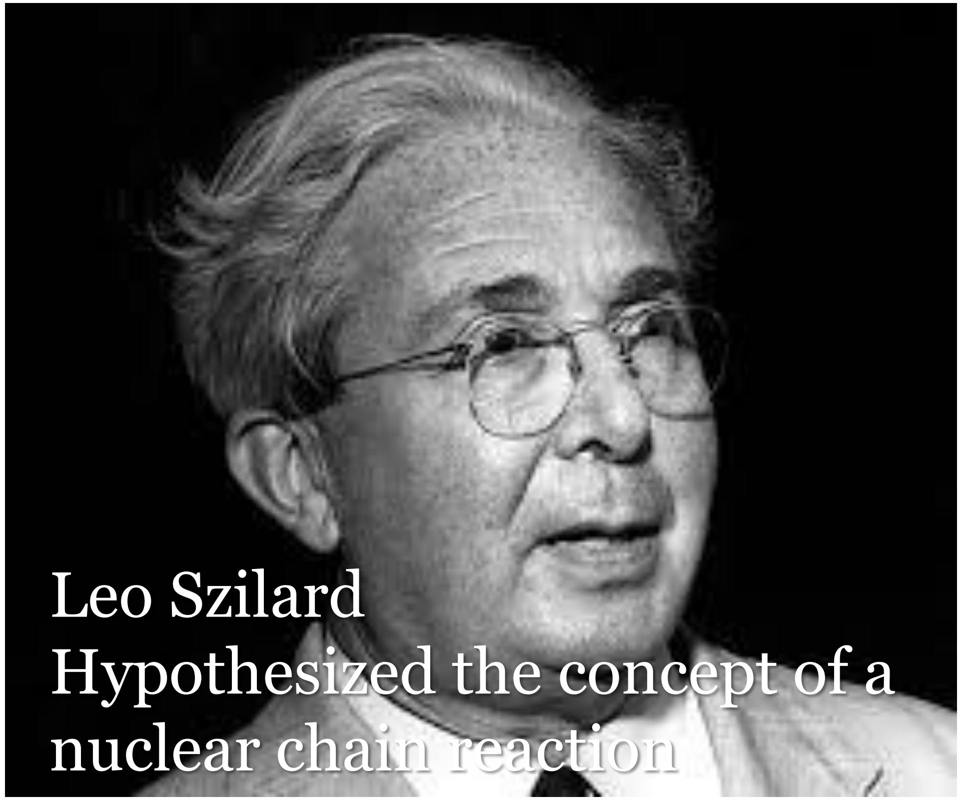

September 12, 1933