Artificial Intelligence & Quantum Computing

Universiteit Leiden
The Netherlands

Aske Plaat
18 september 2018

Outline

- Why is there so much interest in AI?
- What is AI?
- What is QML?
- Will QML be the next disruptor?

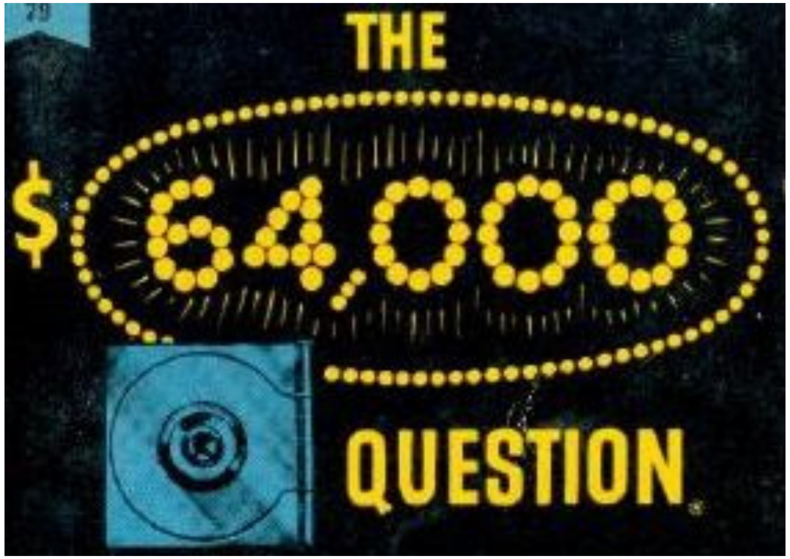

Setting the scene

Why is there so much interest in AI?

Buying Books

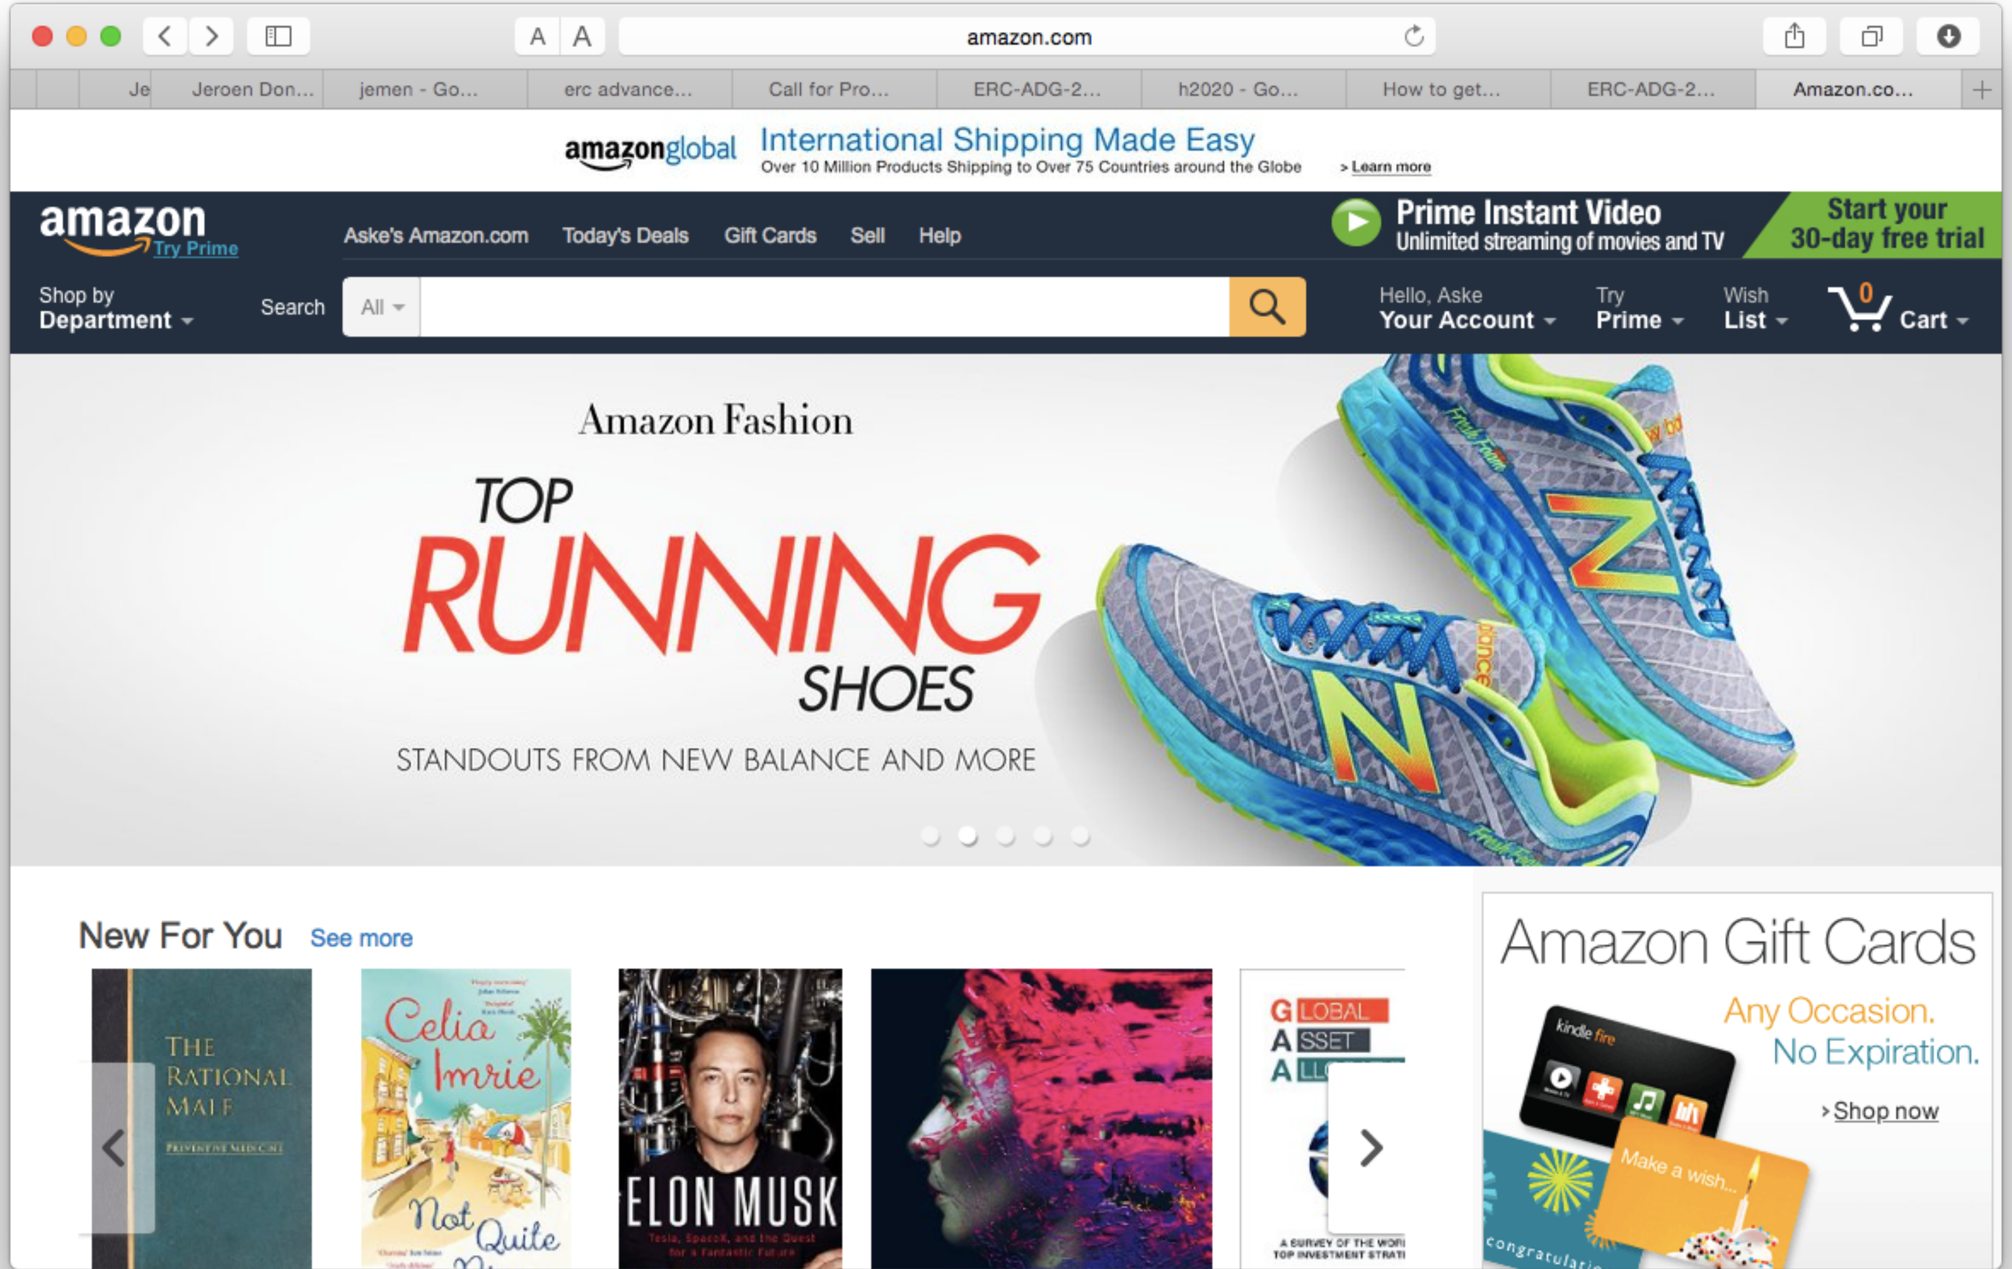

Watching Movies

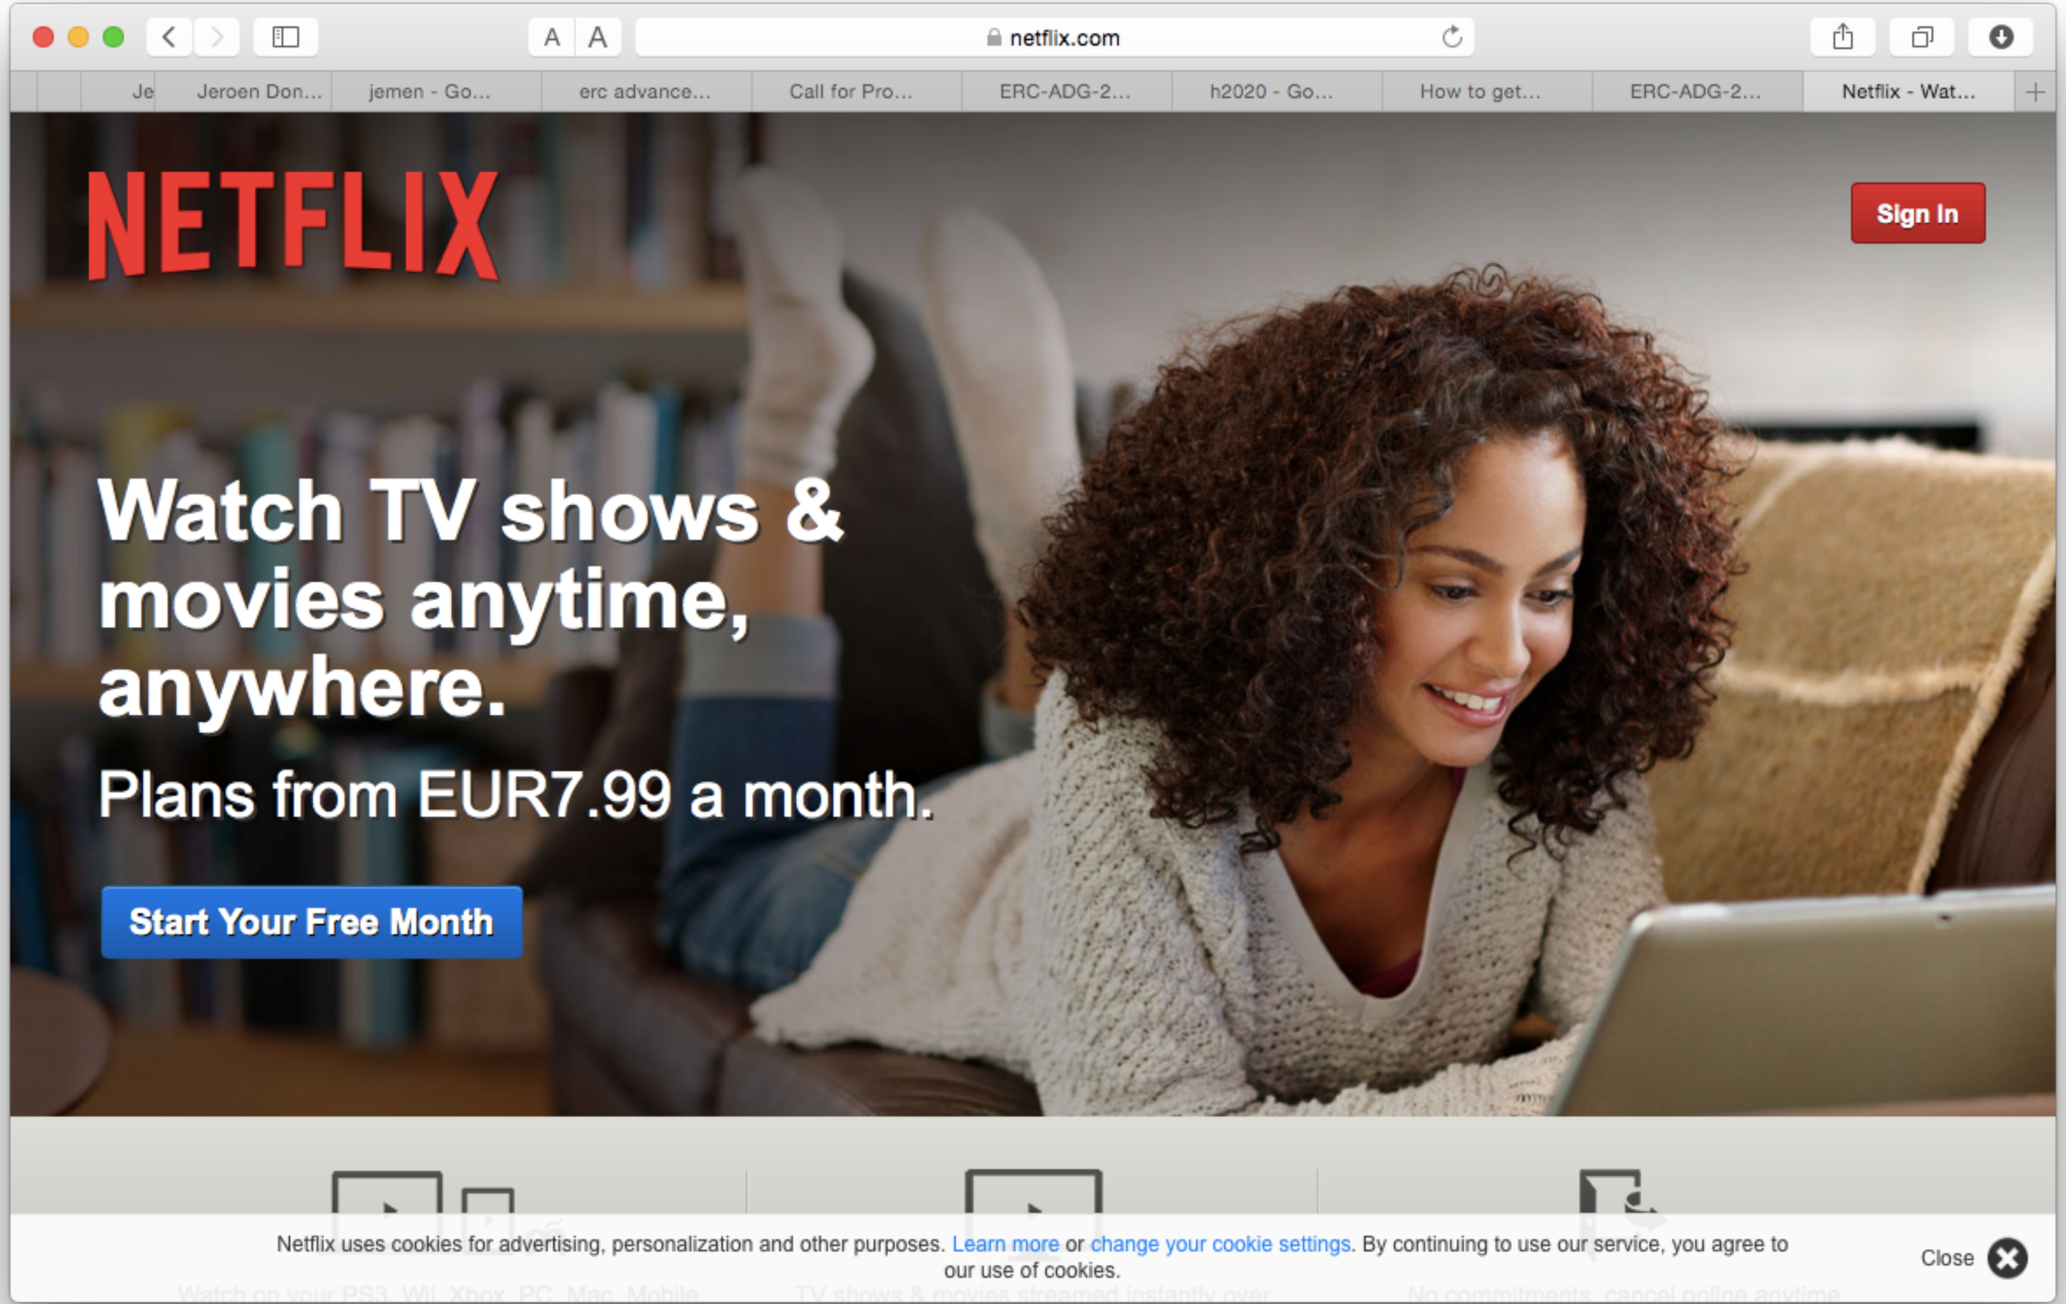

Surfing the Web

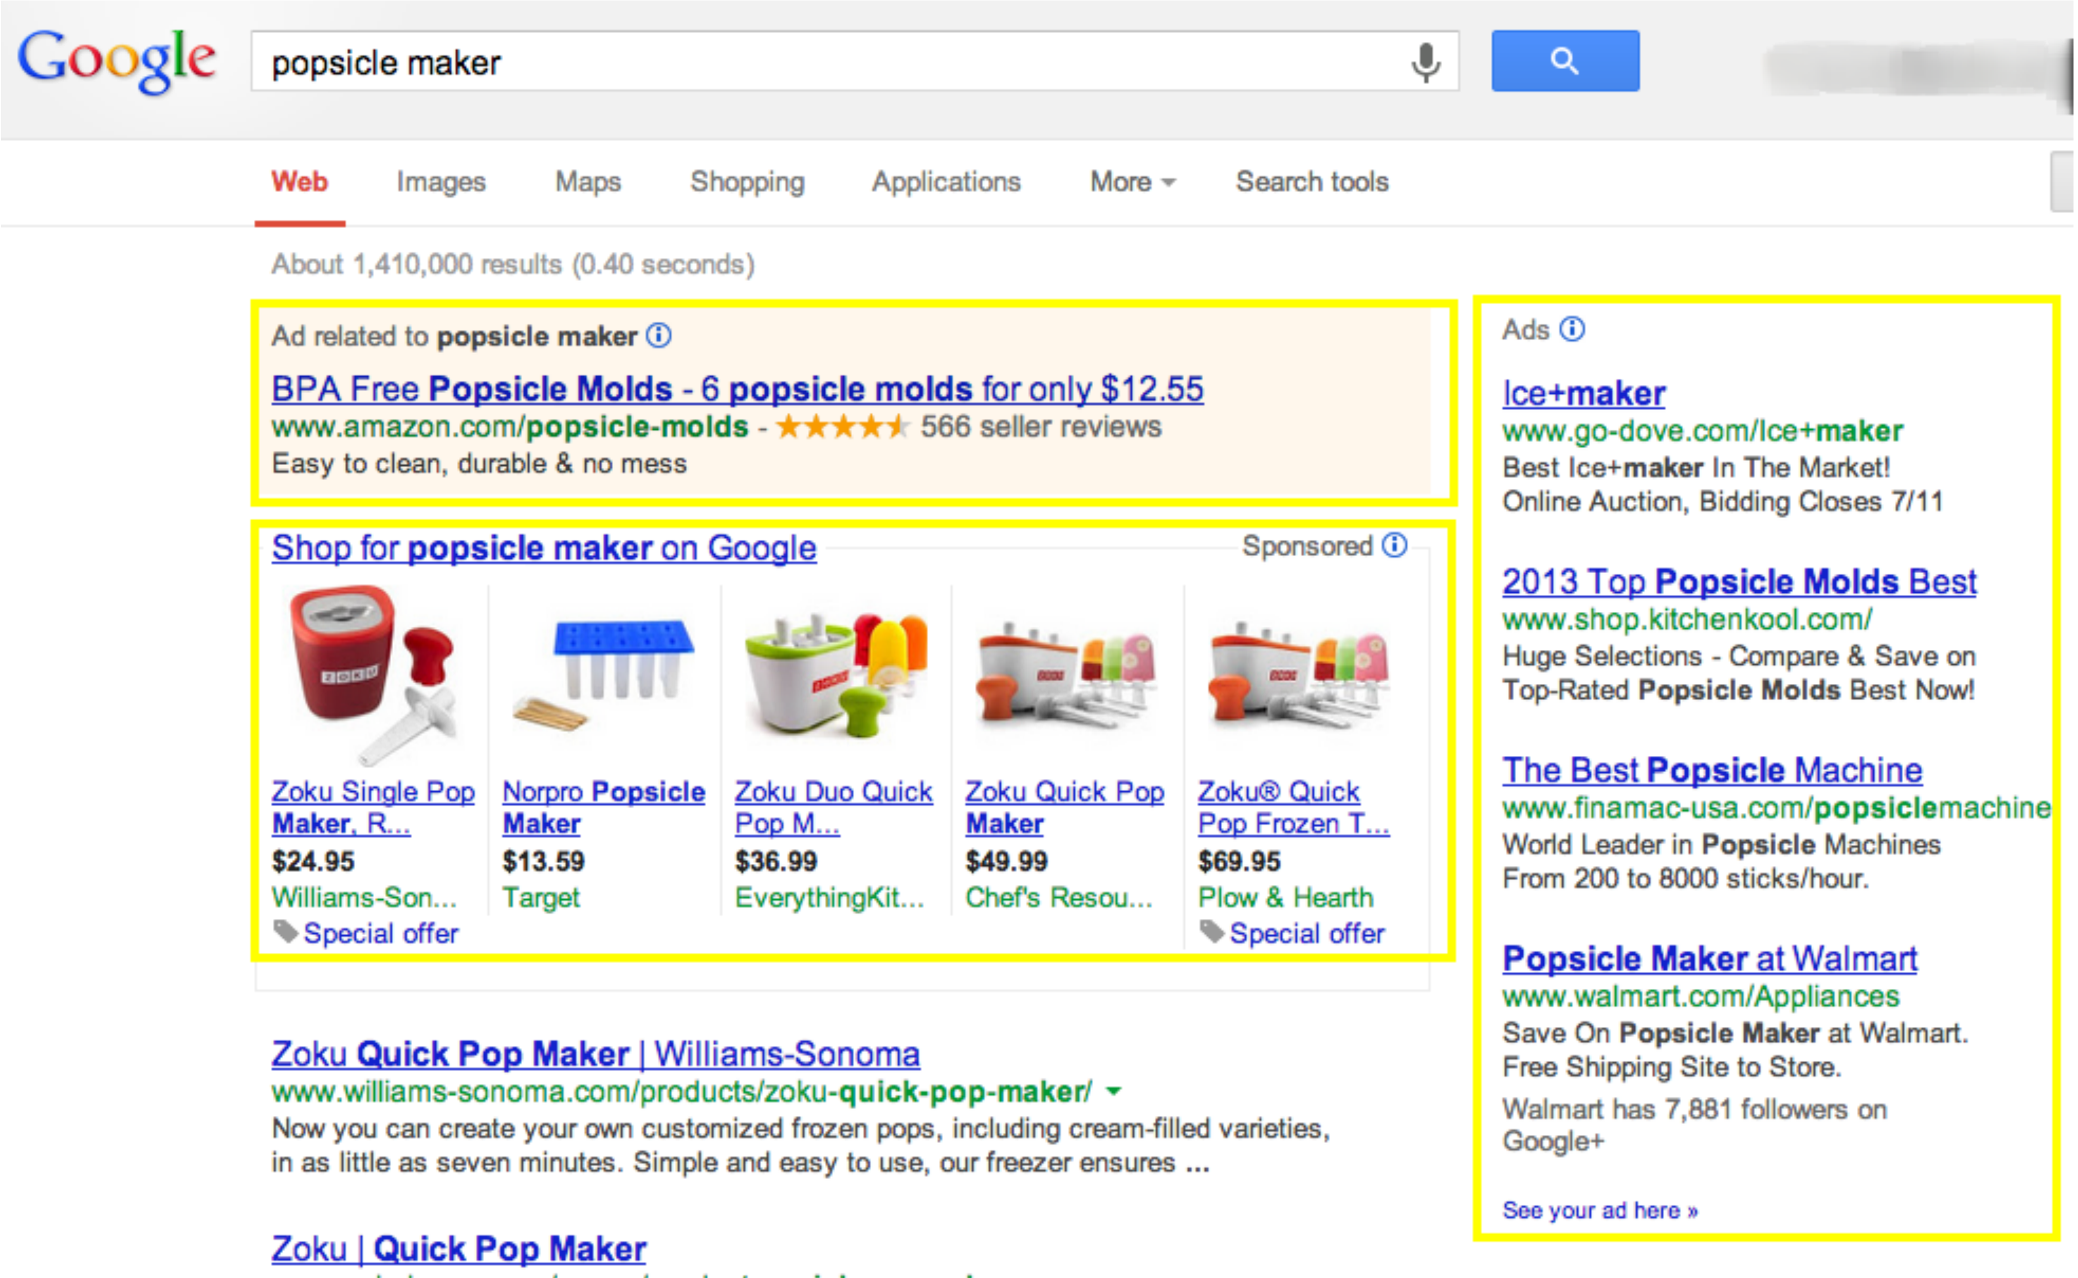

Pattern Recognition

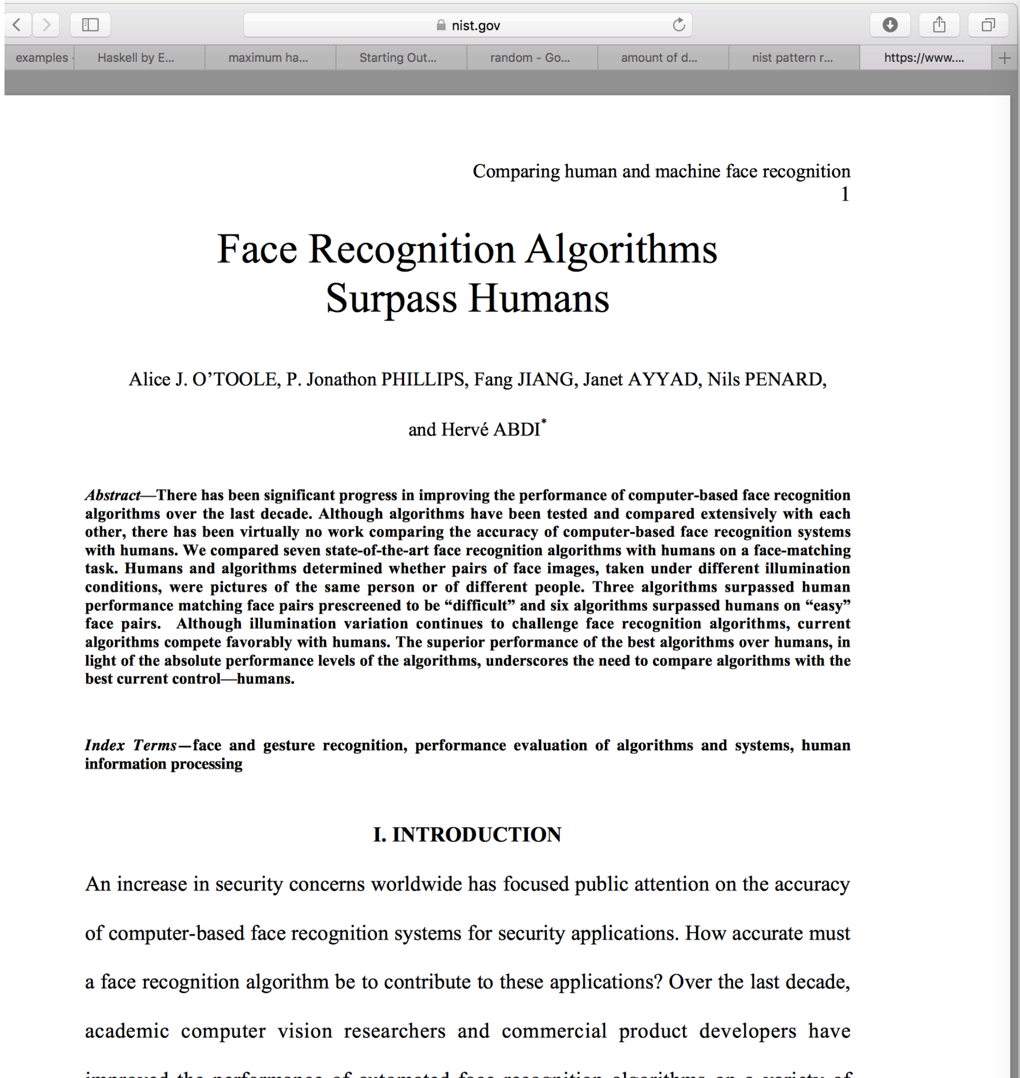

Comparing human and machine face recognition
1

Face Recognition Algorithms
Surpass Humans

Alice J. O'TOOLE, P. Jonathon PHILLIPS, Fang JIANG, Janet AYYAD, Nils PENARD,

and Hervé ABDI*

***Abstract*—There has been significant progress in improving the performance of computer-based face recognition algorithms over the last decade. Although algorithms have been tested and compared extensively with each other, there has been virtually no work comparing the accuracy of computer-based face recognition systems with humans. We compared seven state-of-the-art face recognition algorithms with humans on a face-matching task. Humans and algorithms determined whether pairs of face images, taken under different illumination conditions, were pictures of the same person or of different people. Three algorithms surpassed human performance matching face pairs prescreened to be "difficult" and six algorithms surpassed humans on "easy" face pairs. Although illumination variation continues to challenge face recognition algorithms, current algorithms compete favorably with humans. The superior performance of the best algorithms over humans, in light of the absolute performance levels of the algorithms, underscores the need to compare algorithms with the best current control—humans.**

***Index Terms*—face and gesture recognition, performance evaluation of algorithms and systems, human information processing**

I. INTRODUCTION

An increase in security concerns worldwide has focused public attention on the accuracy of computer-based face recognition systems for security applications. How accurate must a face recognition algorithm be to contribute to these applications? Over the last decade, academic computer vision researchers and commercial product developers have

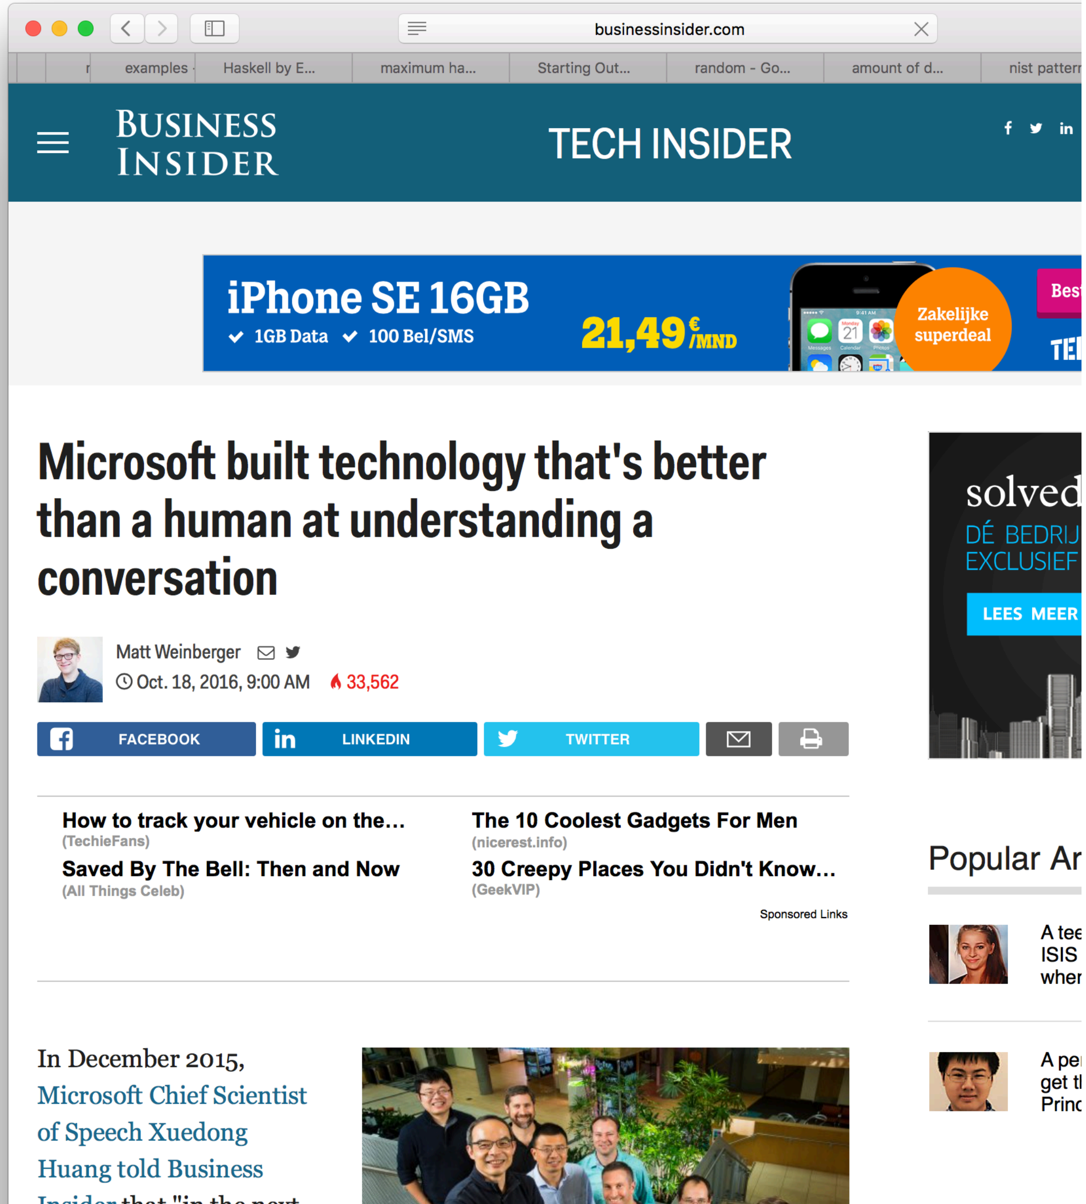

Medical Diagnosis

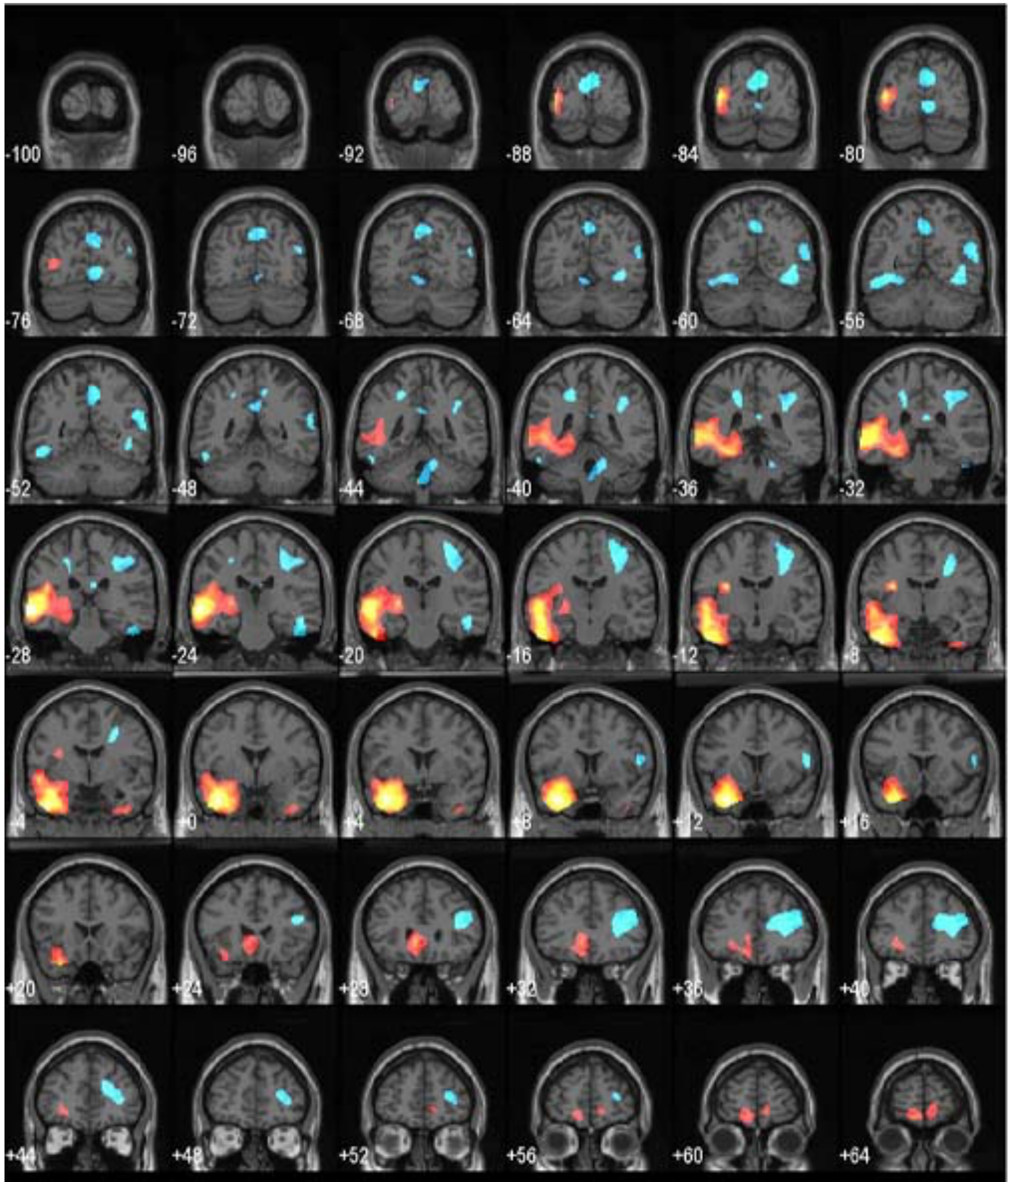

Intelligent Assistants

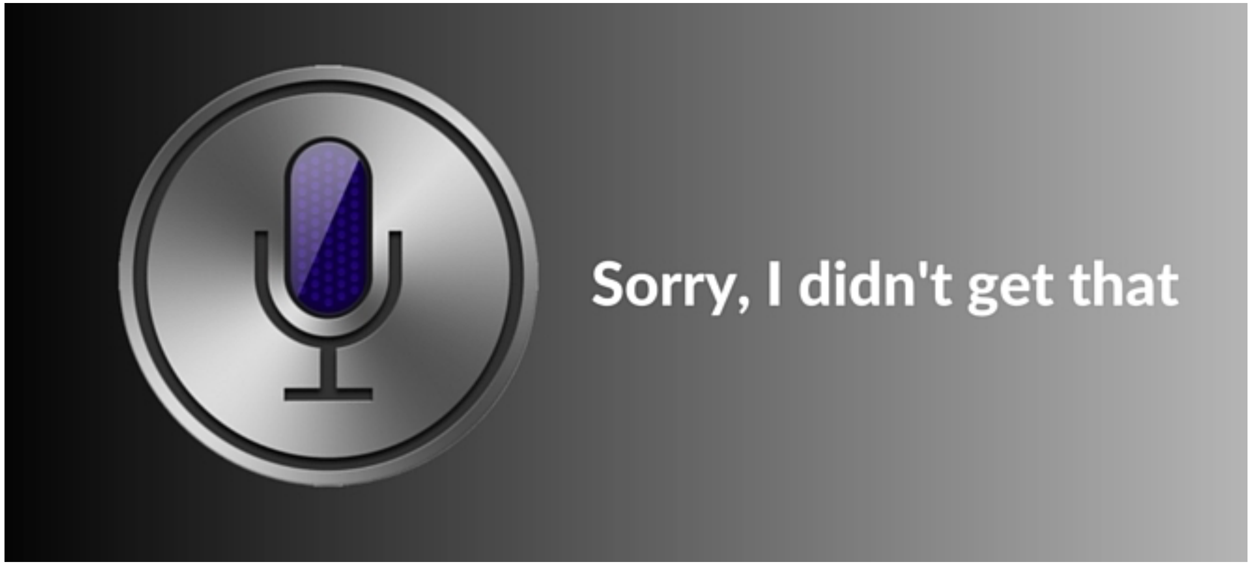

Self-driving cars

Game Playing

Google
amazon
Alibaba.com
f
腾讯
Tencent
ebY
DeepMind
D E Shaw & Co
新浪微博
weibo.com
NATIONAL SECURITY AGENCY
UNITED STATES OF AMERICA
Booking
IBM
Microsoft
Bai du 百度
NETFLIX


The Economist
JUNE 25TH–JULY 1ST 2016
Inside China's Ministry of Truth
Trump in trouble
Who are the Niger Delta Avengers?
The flaws in executive pay
Motorcycles that fly
March of the machines
A SPECIAL REPORT ON ARTIFICIAL INTELLIGENCE


DOUBLE ISSUE
DECEMBER 19, 2016
Puppet of the Year
TIME
DONALD TRUMP
TRAITOR IN CHIEF
PIN HEAD@PIERCEDSKULL
time.com

Disruption

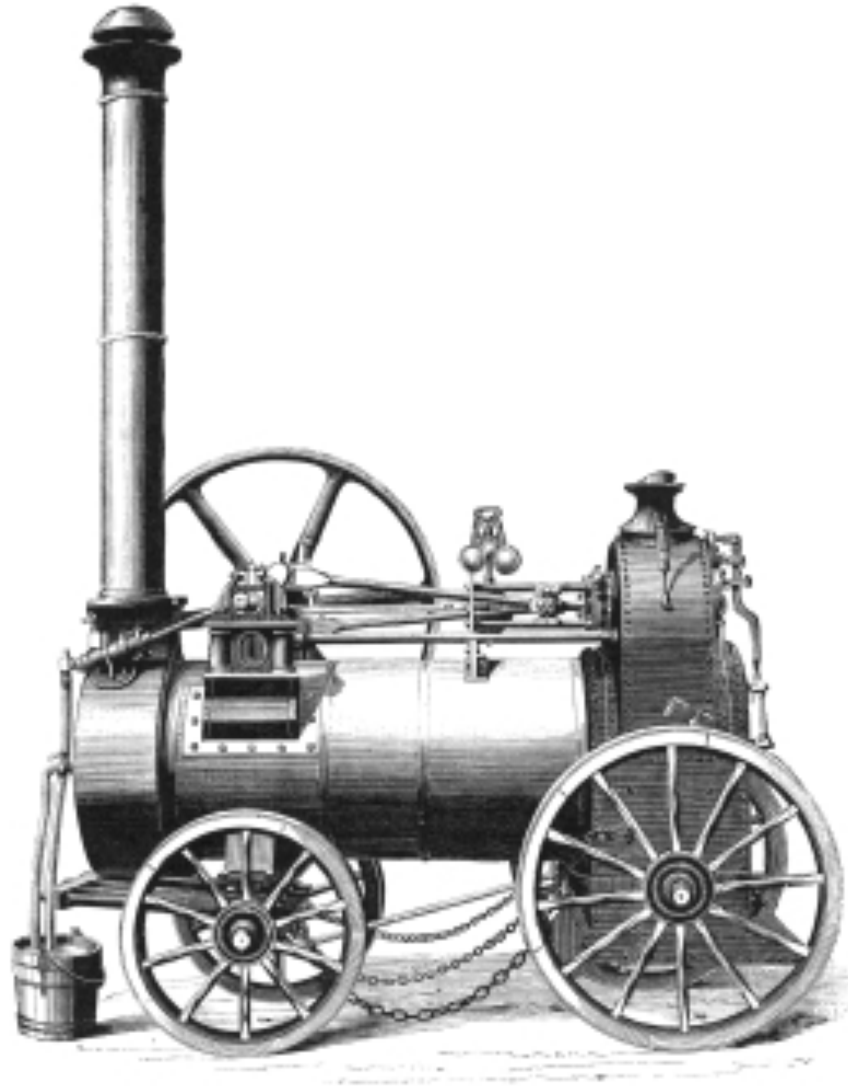

What is Artificial Intelligence?

Is Deep Blue Intelligent?

To be intelligent…

- problem solving
- memory
- creativity
- emotion
- consciousness
- self-awareness

To ***be***...

- problem solving
- memory
- creativity
- emotion
- consciousness
- self-awareness

To *act*…

The Turing Test

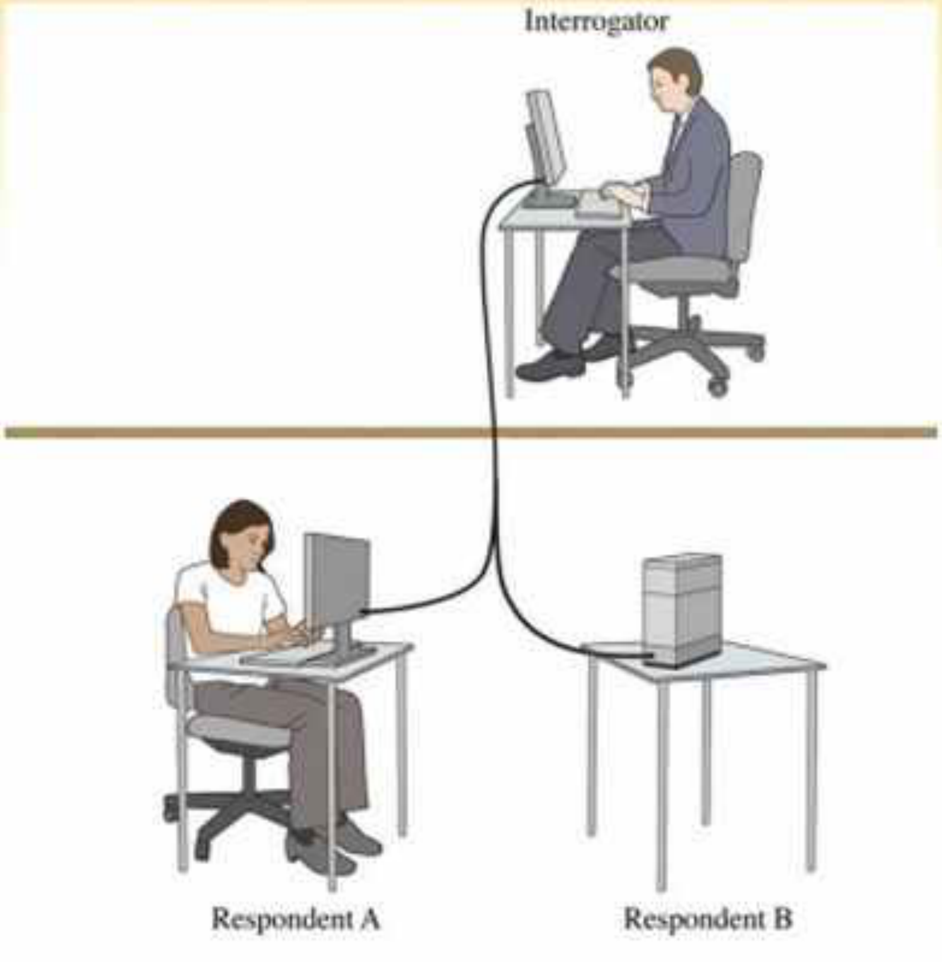

Figure 13.2
In a Turing test, the interrogator must determine which respondent is the computer and which is the human

Acting Intelligent

- sensing
- **adaptation**
- actuating

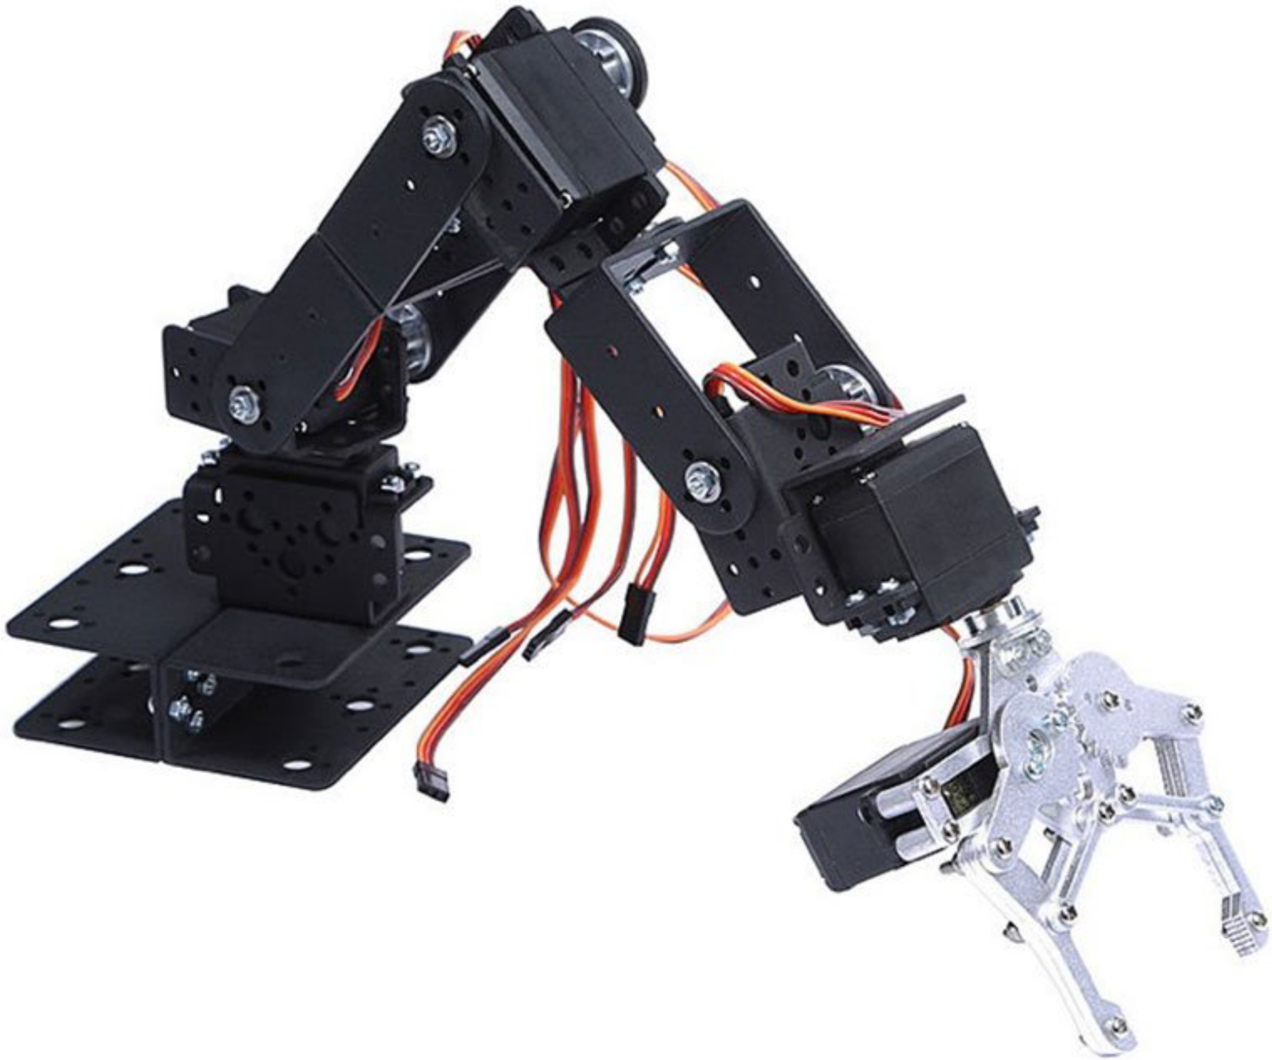

Adaptation

- logical reasoning
- memory
- learning/generalization

Machine Learning

- logical reasoning
- memory
- **learning/generalization**

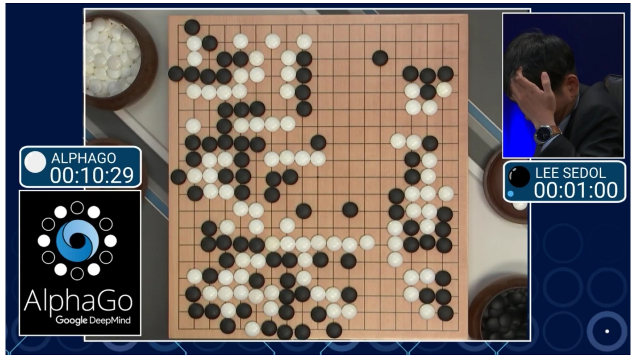

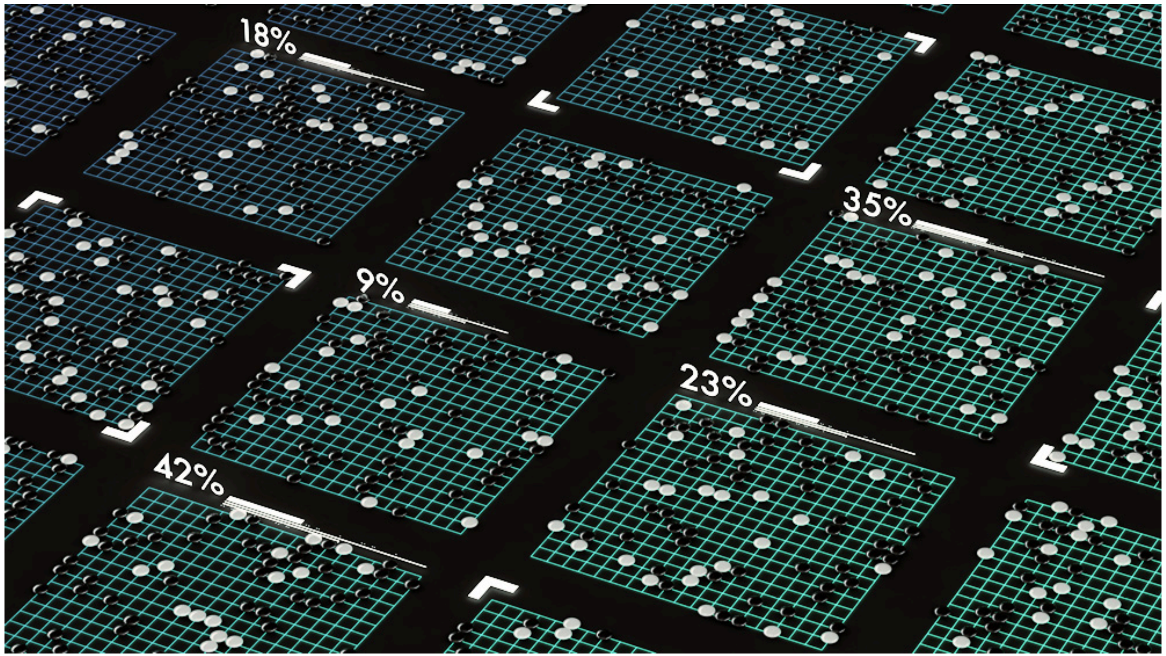

Supervised Learning

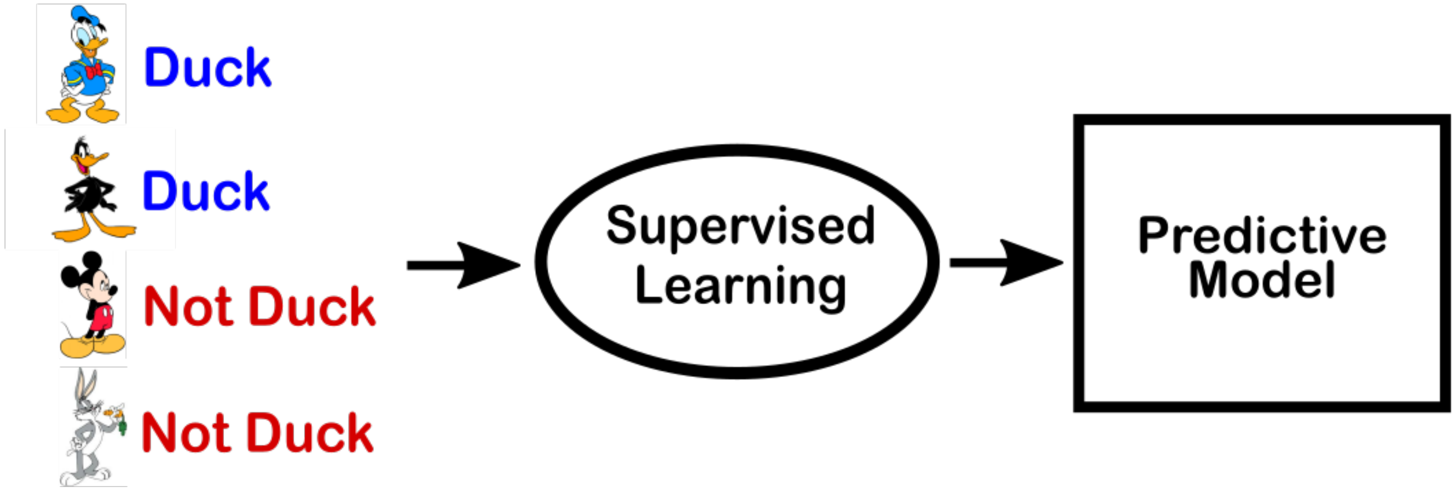

- learning = generalization = approximating a mapping

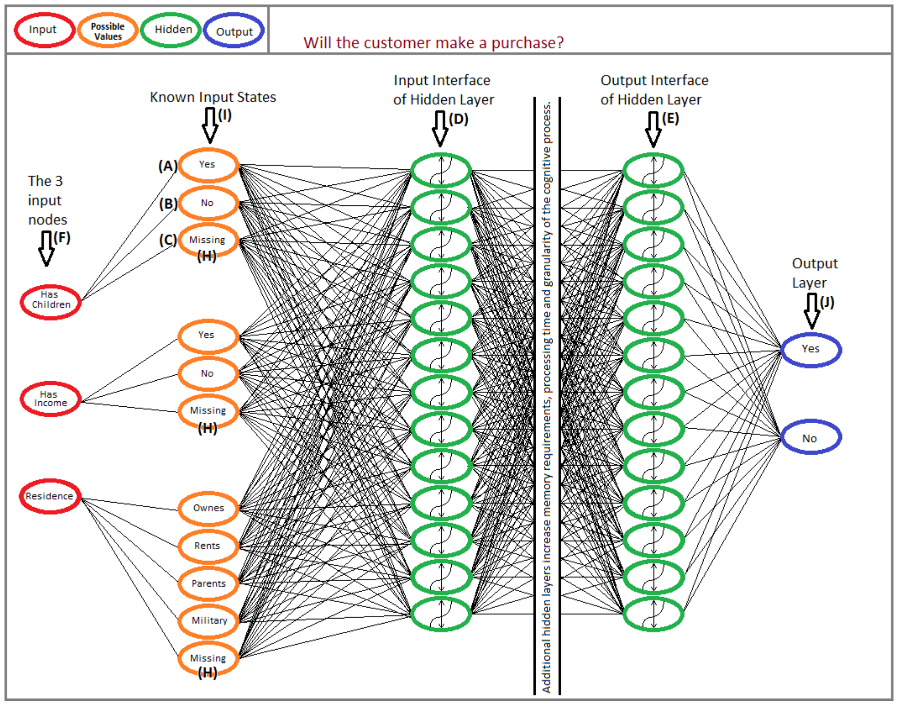

No hard-to-define concepts

- Self awareness
- emotion
- consciousness
- creativity

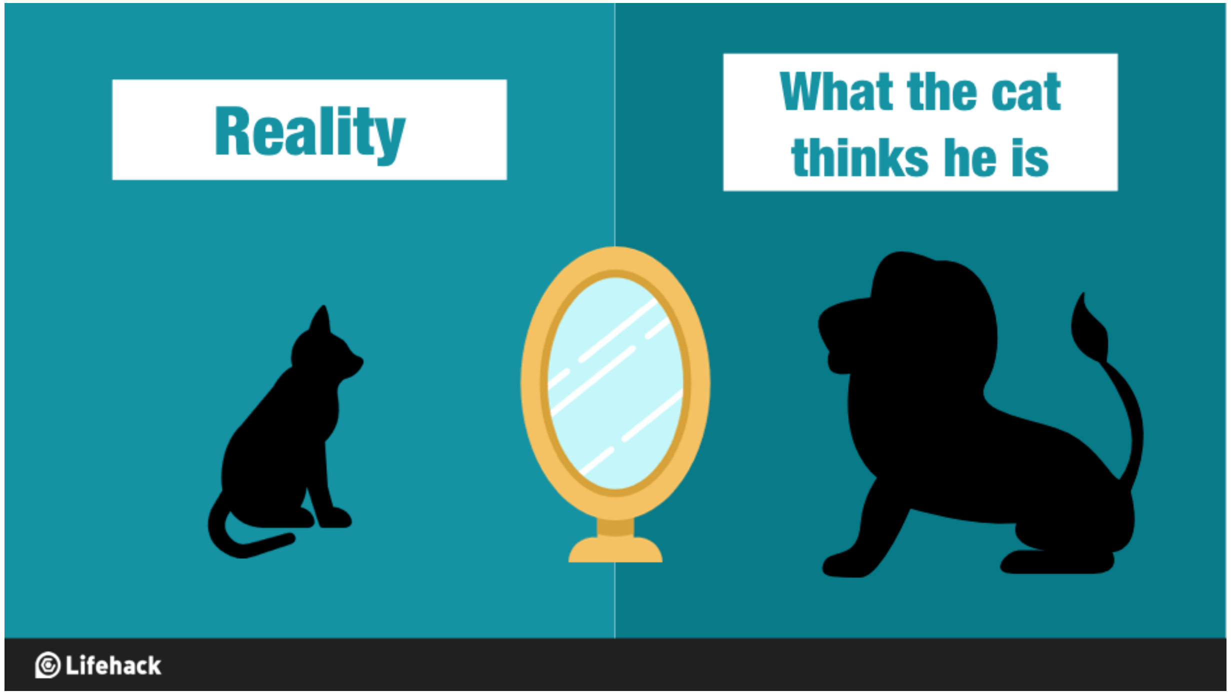

Yes clear problem domains

Yes clear problem domains

What technology is behind the AI breakthroughs?

AI

OpenML

1. Algorithm

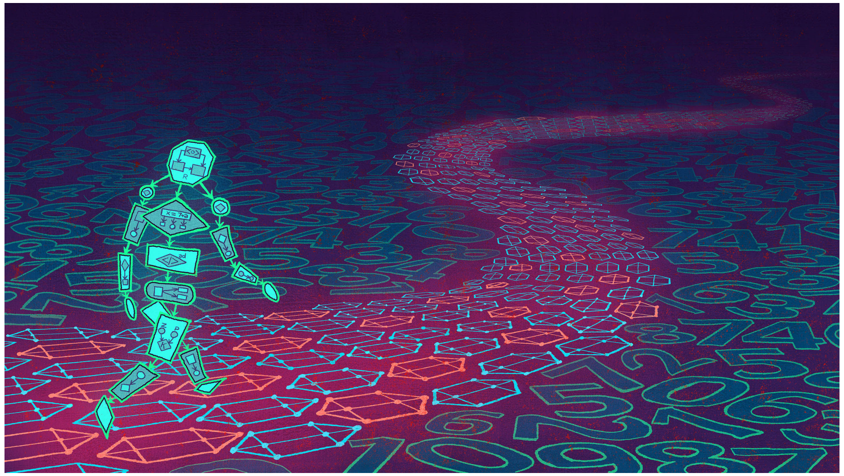

2. Data

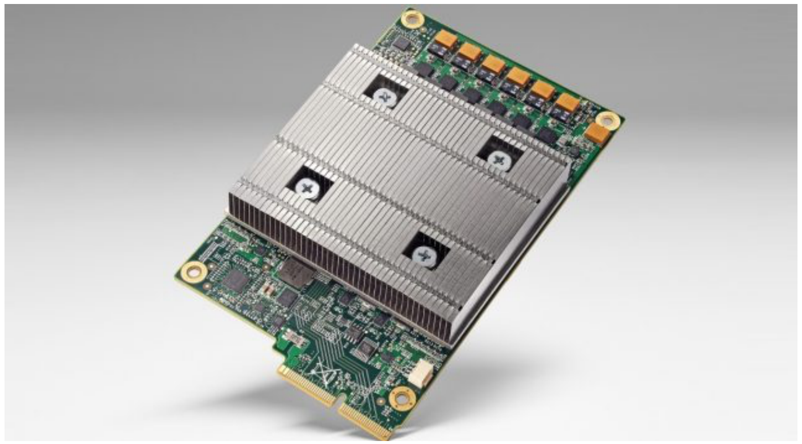

3. Speed

1. Algorithms

Symbolic

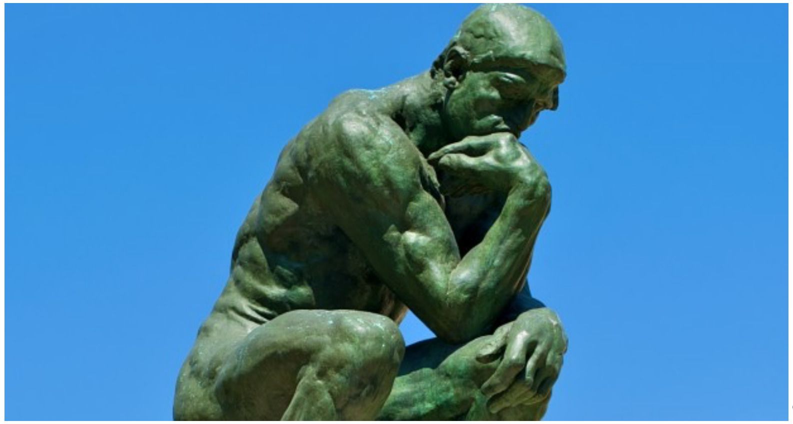

Connectionist

Symbolic

- Philosophy, logic, system 2
- Top down
- Expert systems
- Ontologies, resolution,
- Agents, Planning
- Tabular Q Learning, SVM, PCA
- Markov Decision Problems, Mcts

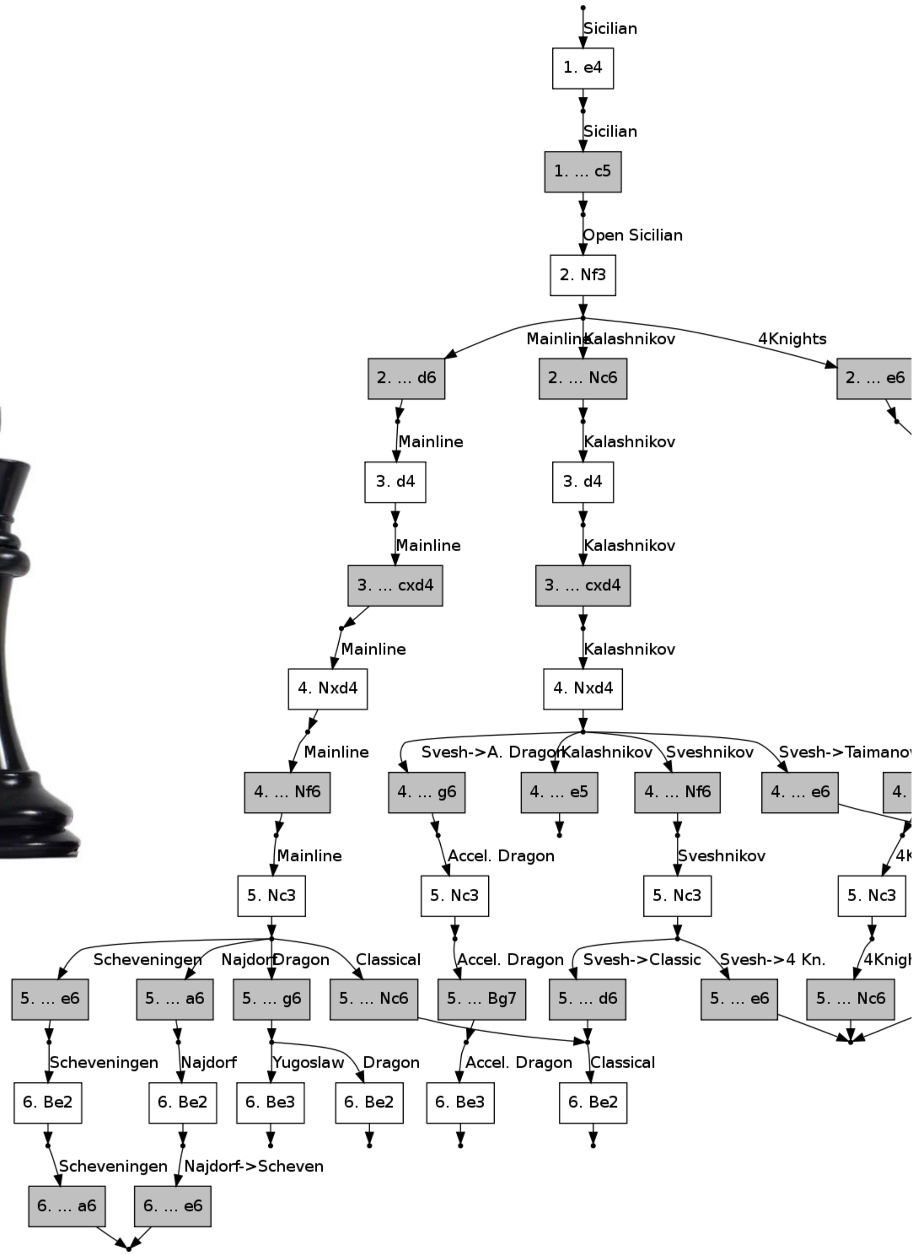

Connectionist

- Biology, neural net, evolutionary, robotics, life, system 1
- Generalization, function approximators
- Perceptrons
- Deep Learning
- Bottom up
- CNN, RNN, LSTM
- Skips Resnet, Regularization
- Batch normalization vanishing gradients

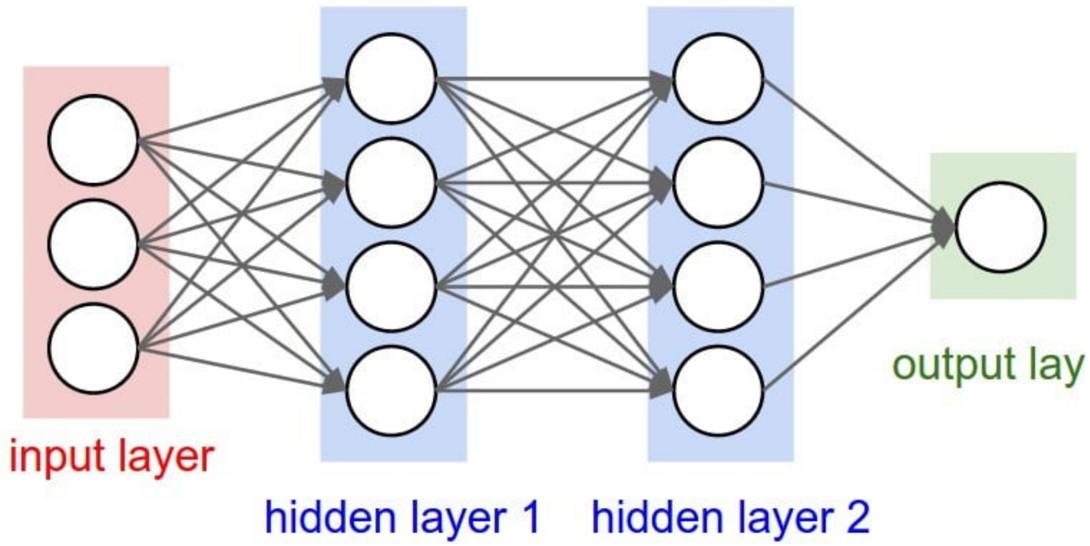

Issues

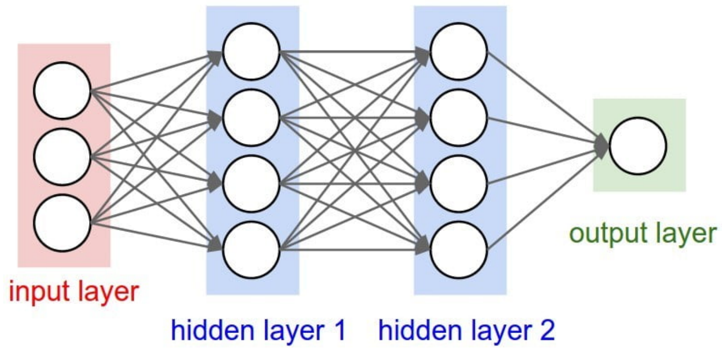

- A net is a function mapping input on output
- In theory a one hidden layer net is a sufficient function approximator
  Many layers work better in practice
- Learning a function is generalizing example states, optimization of underconstrained multidimensional problem. Many more parameters than examples
- Regularization tricks to get convergent learning
- Works great in practice, no satisfactory theory of deep learning

Notable Successes

- Symbolic reasoning: Chess

- Agent simulations: Emergent Social Norms, Negotiation
- Evolutionary Optimization in Industry

- Deep Convolutional Nets: Image recognition

- Recurrent Nets: Speech recognition

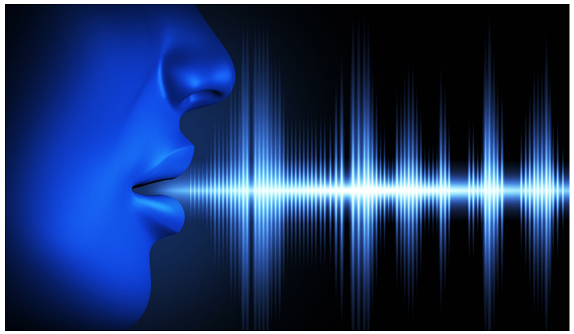

- Search+Net: Go

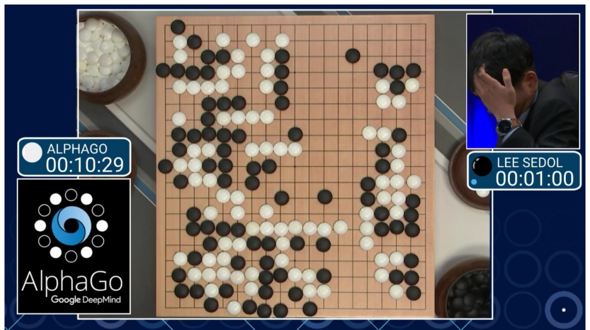

Kahneman

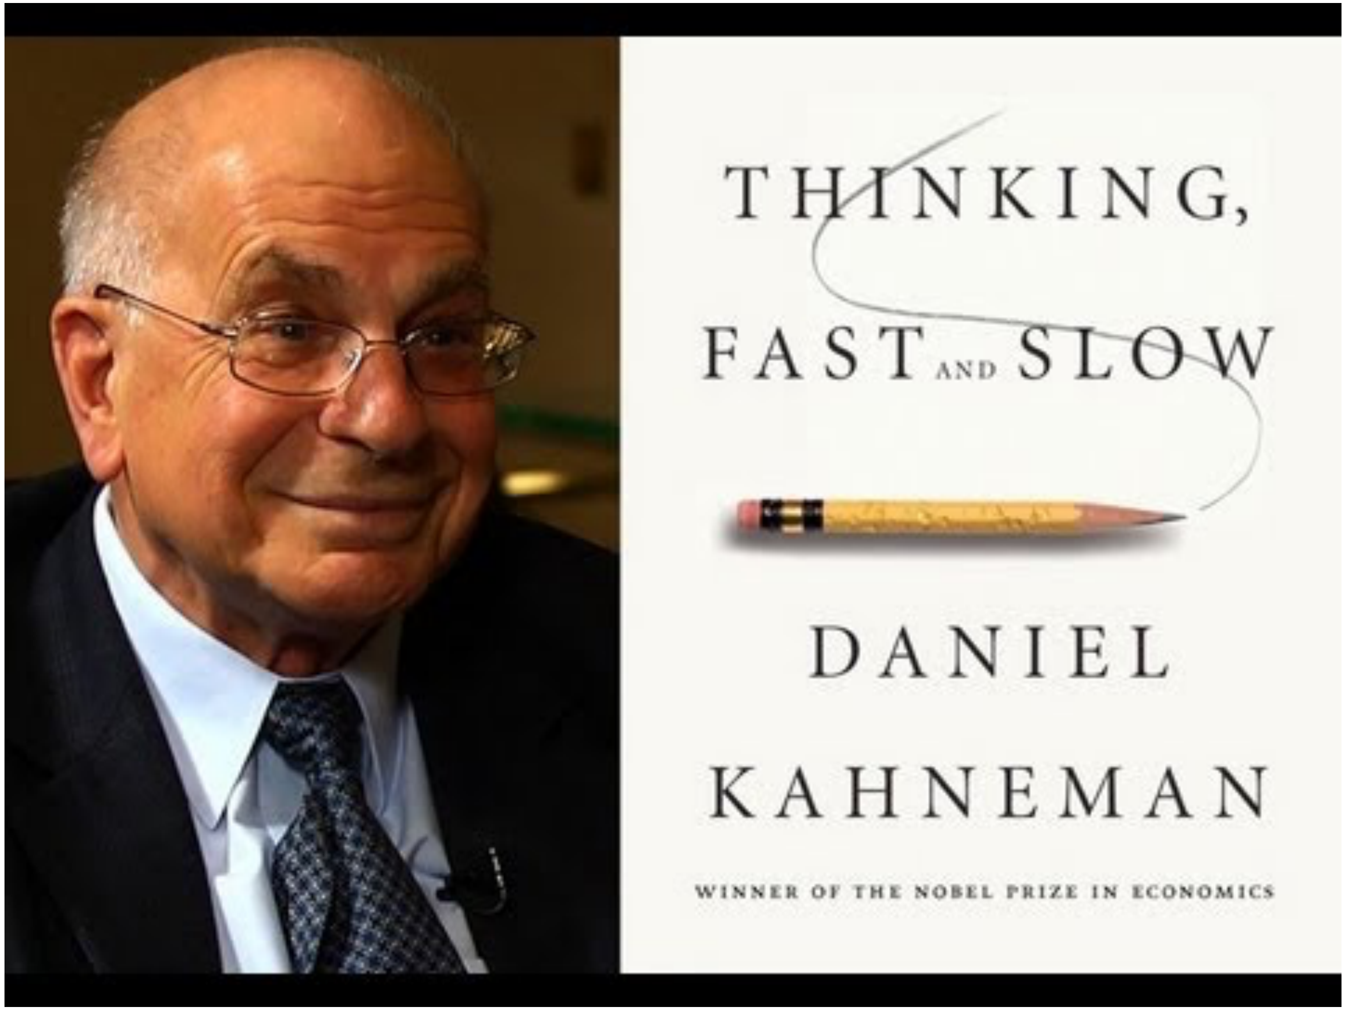

1: reflex. 2: reasoning

2. The Importance of Data for AI

Data

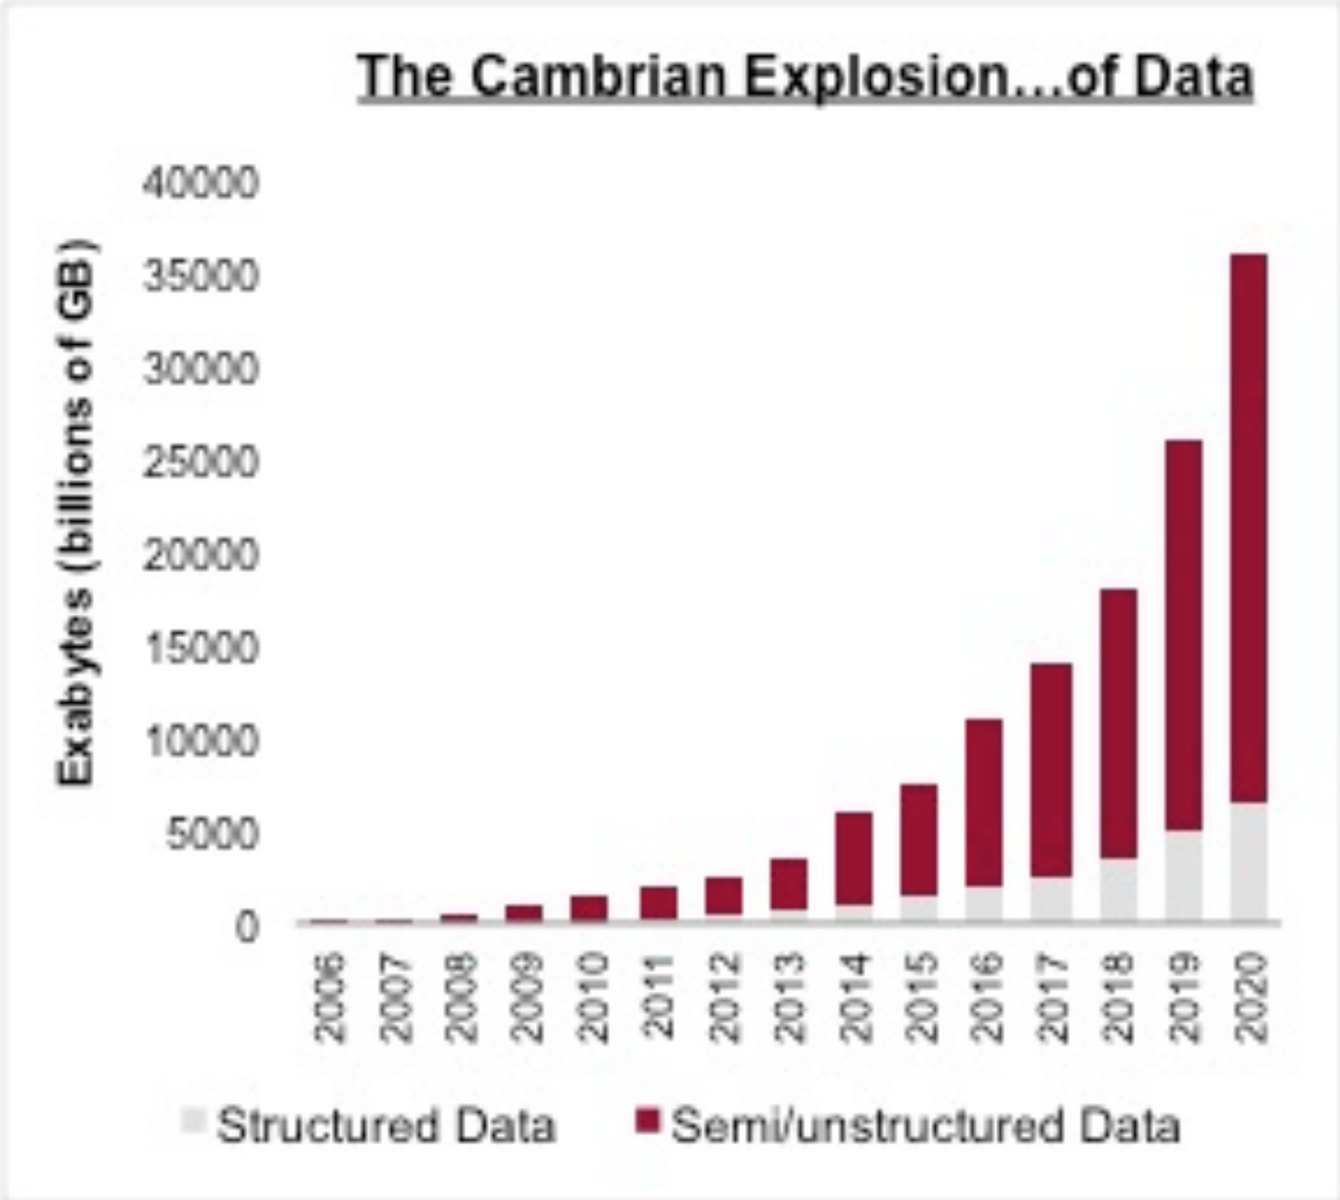
The Cambrian Explosion…of Data
Exabytes (billions of GB)
40000
35000
30000
25000
20000
15000
10000
5000
0
2006
2007
2008
2009
2010
2011
2012
2013
2014
2015
2016
2017
2018
2019
2020
Structured Data
Semi/unstructured Data

Function Approximator
Generic Mapper

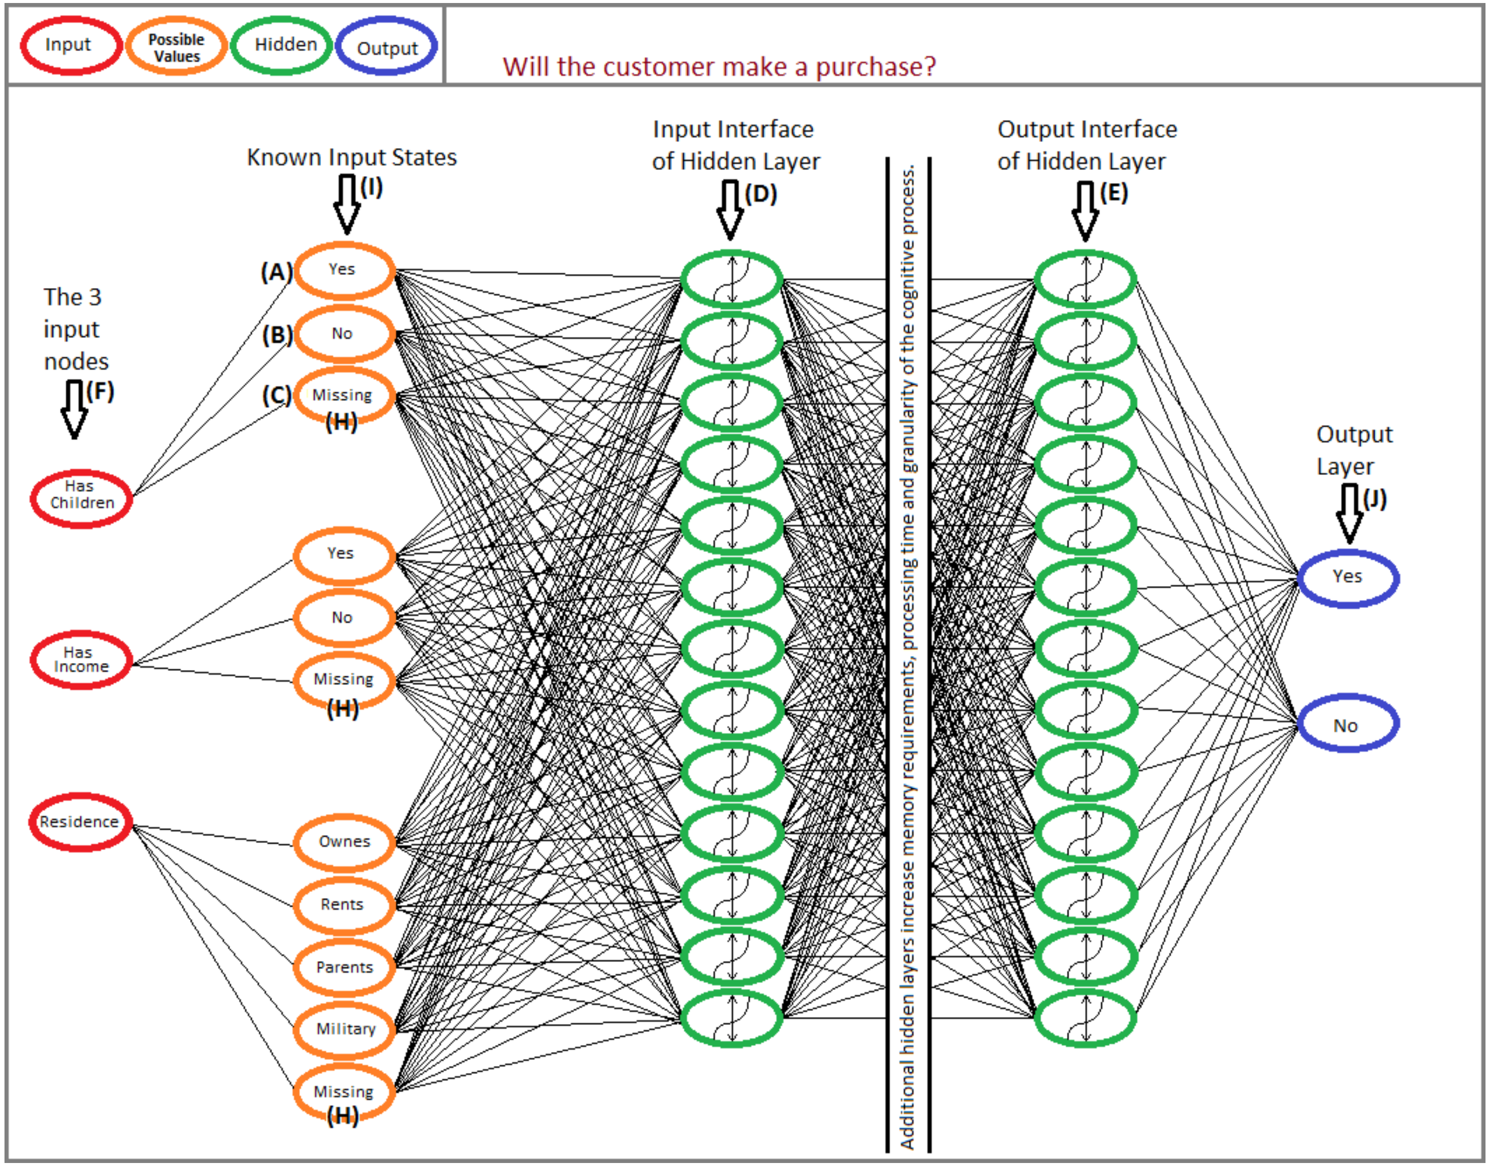

Supervised Learning

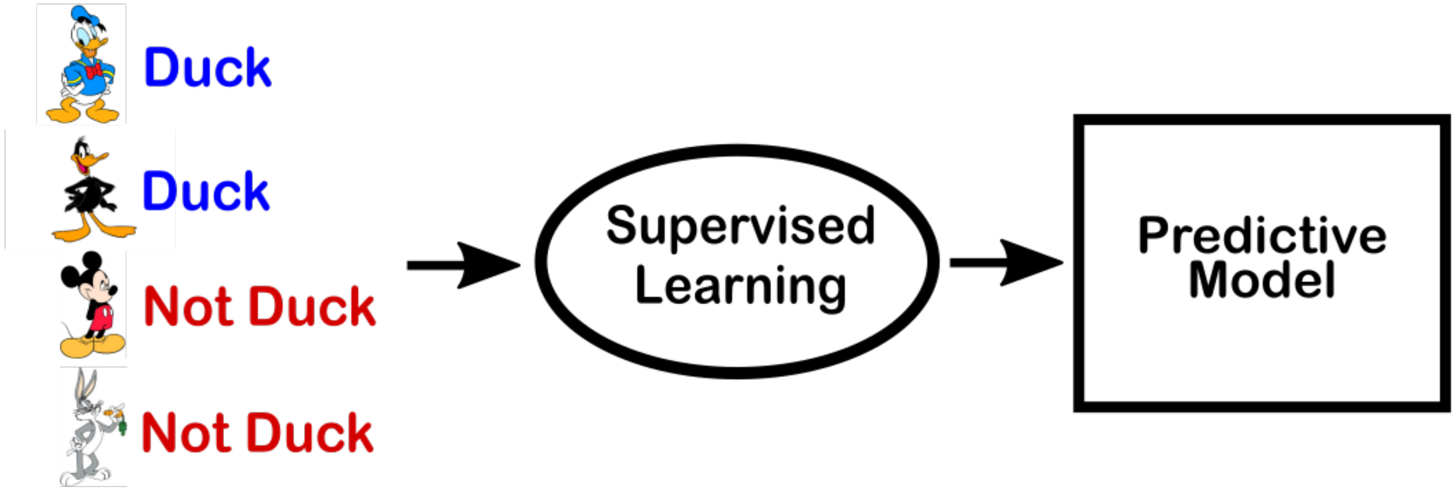

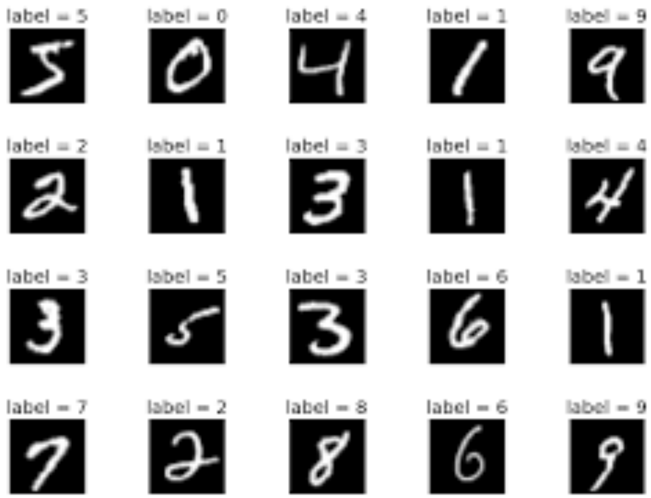

Reinforcement Learning

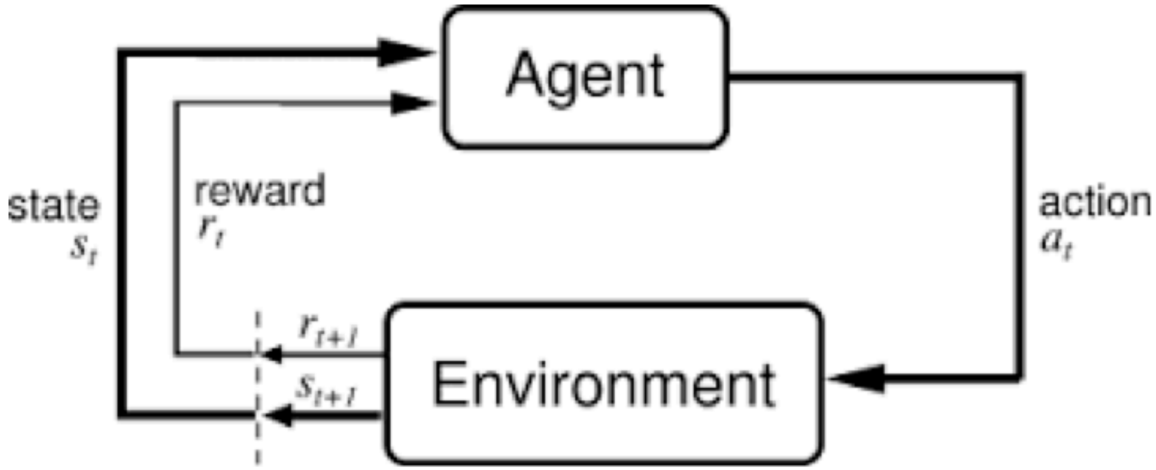

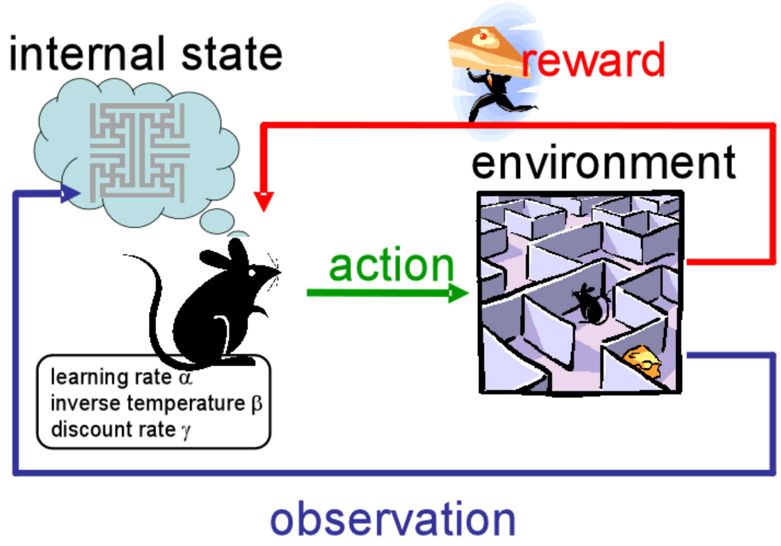

- Internal states have no reward, only actions leading to other states
  Rewards must be derived

- symbolic
  chess: alphabeta search + eval of hand coded features

- connectionist
  backgammon: NN of hand coded features
  atari: Deep Q-learning features DQN

- symbolic + connectionist
  go: MCTS search generating training examples for DQN

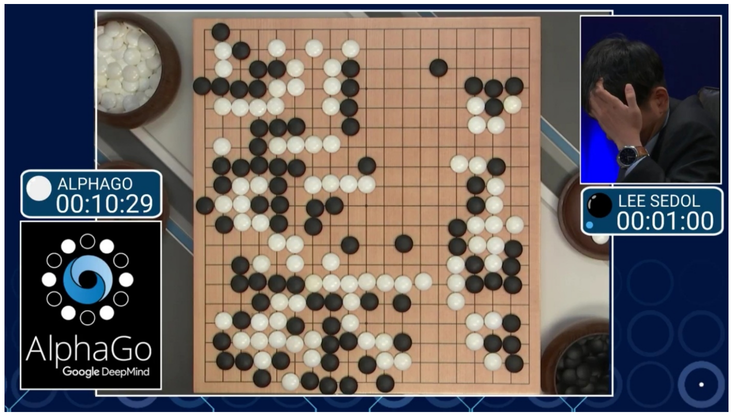

Supervised / Reinforcement

- Supervised Learning: EXPLICIT big data
- Reinforcement Learning: IMPLICIT big data

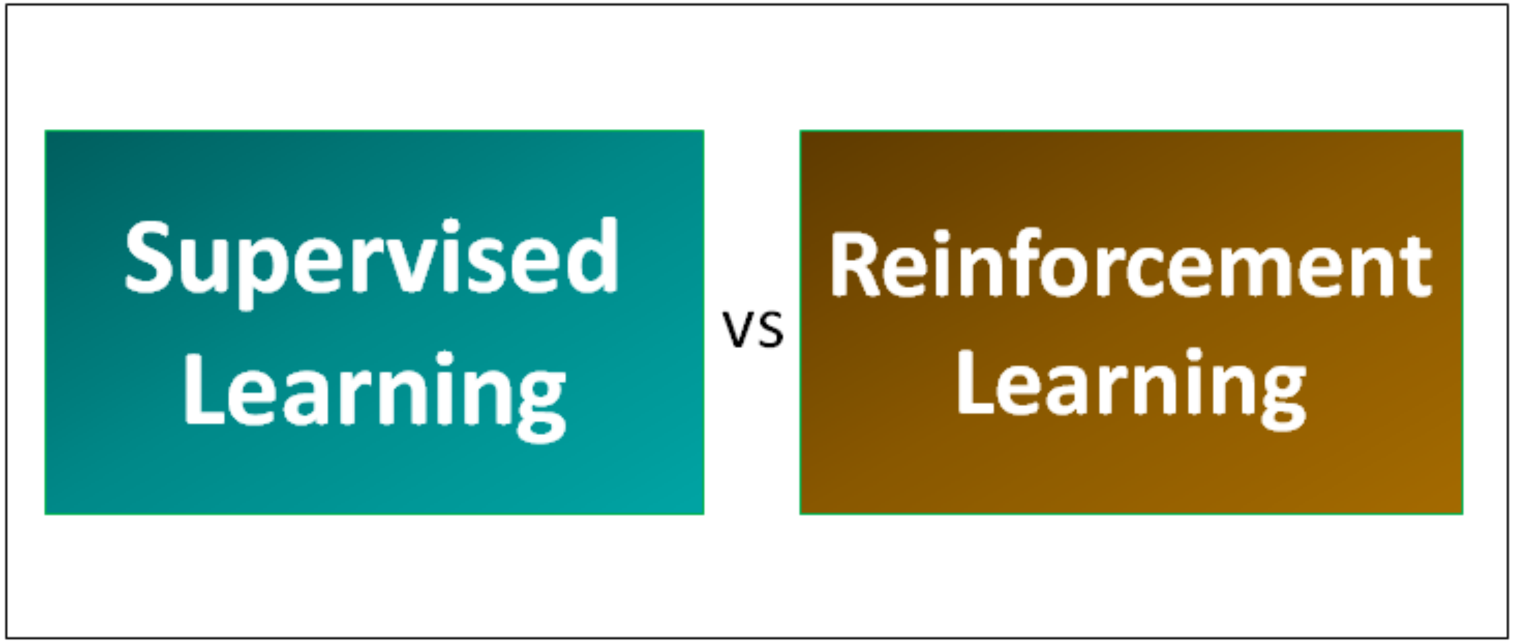

Data

- **Supervised learning** big labelled data sets
- Consumer
- **Connectionist generalizes** data into theory
- Image, Speech recognition
- Preference "recognition"

- **Reinforcement learning** uses simulators to generate examples
- Producer
- **Symbolic** generates data from theory
- **Planning** in Games, Markov
- Agent Simulators

3. The Importance of Speed for AI

3. Speed

Cost of Computing Power Equal to an iPad 2

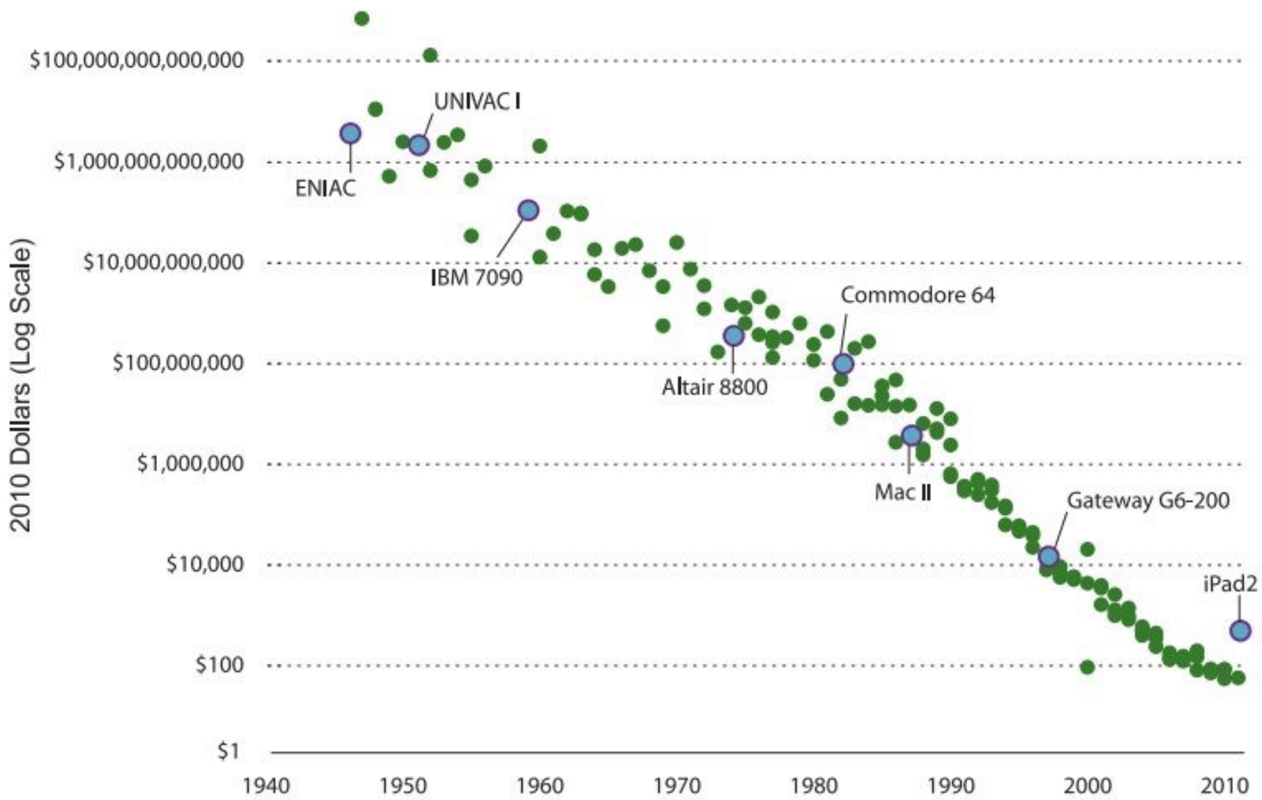

Note: The iPad2 has computing power equal to 1600 million instructions per second (MIPS). Each data point represents the cost of 1600 MIPS of computing power based on the power and price of a specific computing device released that year.

Source: Moravec n.d..

Speed

- Matrix operations : GPU (highly parallel matrix ops)
- Software packages: Theano, Torch, Tensorflow, Python
- Open source

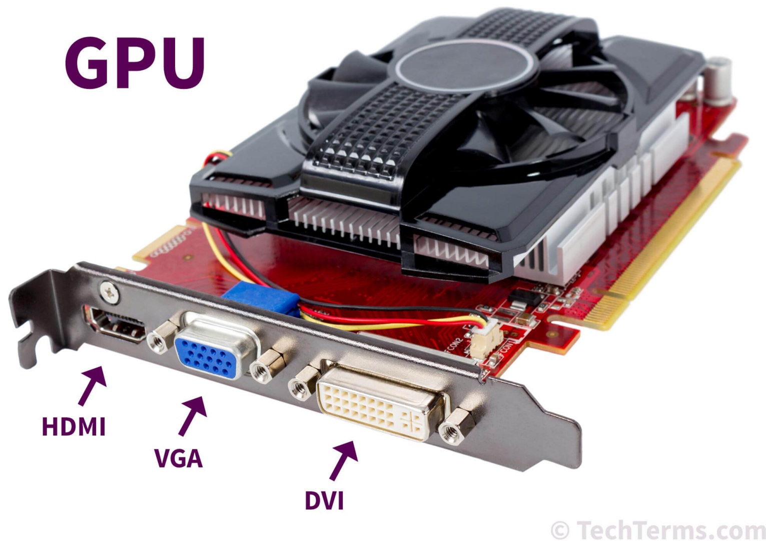

Interlude: Benchmarks

deep supervised learning	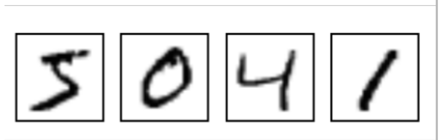	continuous domain	
deep reinforcement learning		imperfect information/multi agent	
MCTS+DQN	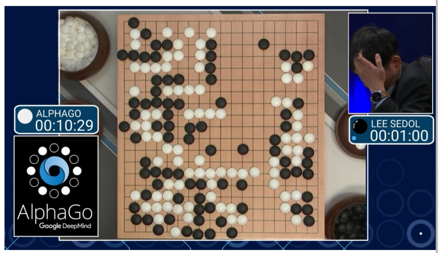	imperfect information	
deep Q-learning	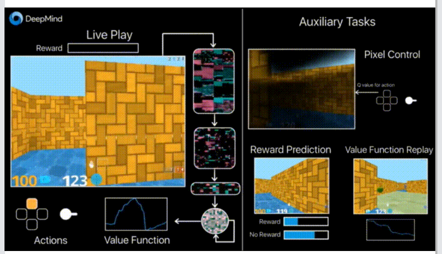	imperfect information/ cooperation	

Summary

1. Algorithms: Symbolic & Connectionist
2. Data: Supervised explicit & Reinforcement implicit
3. Speed: Matrix operations on GPUs & Open Source

Need for Speed

- If intelligence is learning to solve large combinatorial problems

- Then Machine Learning converts compute power into intelligence

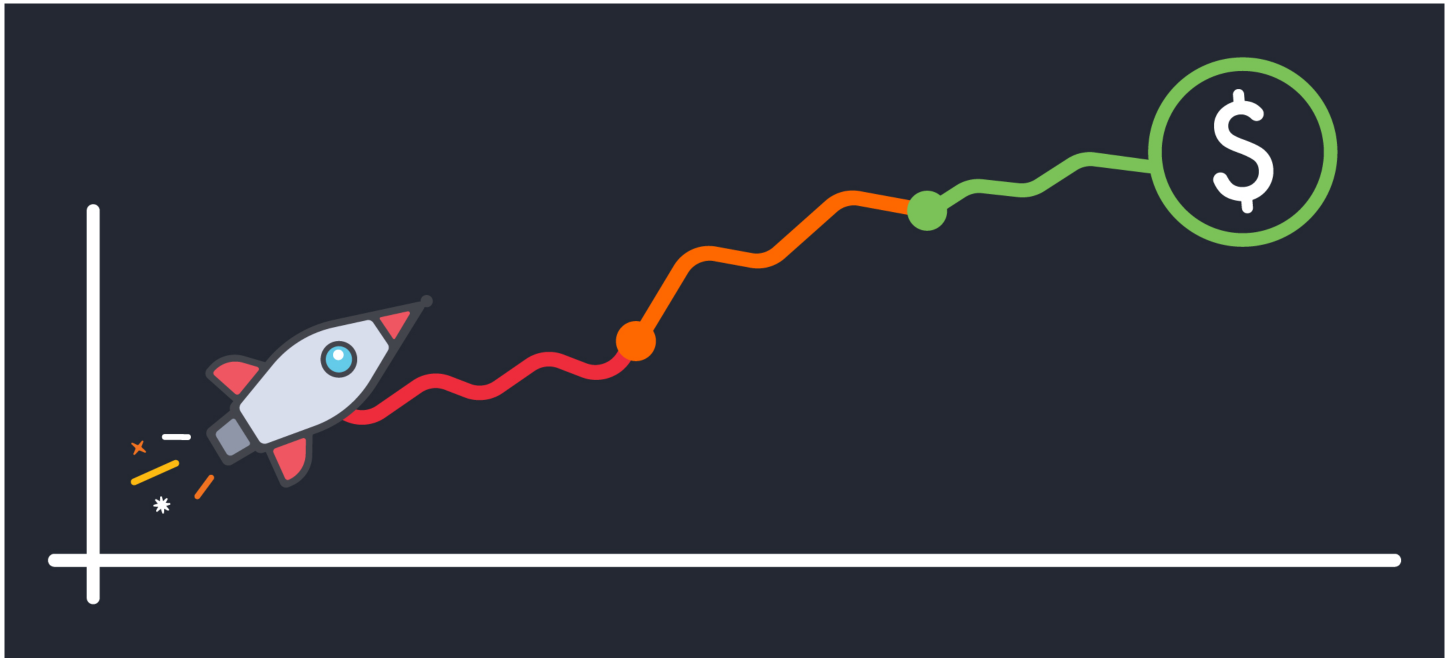

Future AI?

- Present AI
  - Approximated some complex problems better than humans.
    Image & Speech Recognition, Recommendation, Game playing
- Future AI
  - Autonomous systems: cars, Search and rescue, Space exploration
  - Medical diagnosis, Algorithmic trading
  - Profiling: security, insurance
  - Efficient logistics, humanitarian aid
  - Efficient, sustainable production

Challenges in ML

- Domains with:
  - Sparse inputs
  - Imperfect information
  - Delayed credit assignment
  - Large branching factors
- Learning:
  - Transfer learning/ model lifting
  - Imitation
  - Life long learning
  - Hierarchical
  - Metalearning
  - Robust/1 pixel problem

Challenges in ML

- Domains with:
  - Sparse inputs
  - Delayed credit assignment
  - Large branching factors

- Learning:
  - Transfer learning/ model lifting
  - Life long
  - Hierarchical
  - Metalearning
  - Robust/1 pixel problem

If we have all the compute power in the world,
Will these challenges be solved?

What is QML?

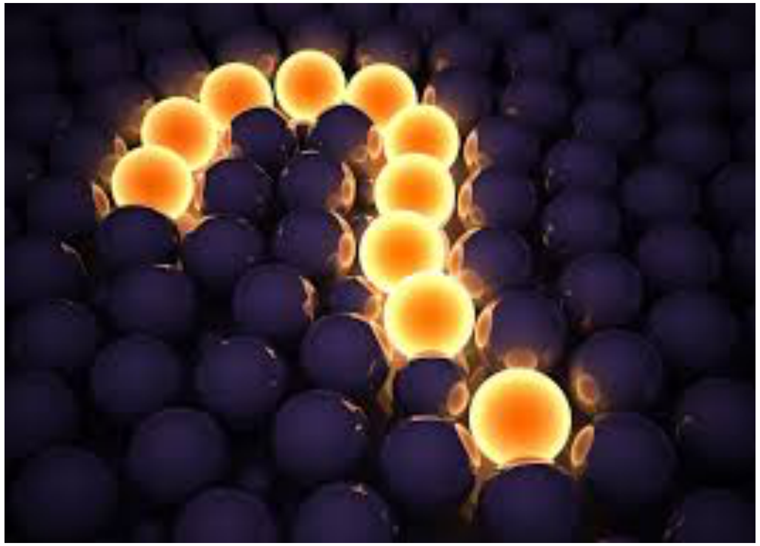

1. Algorithms

- Annealing

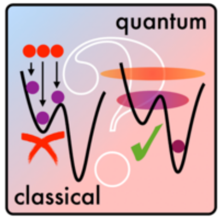

- Superposition

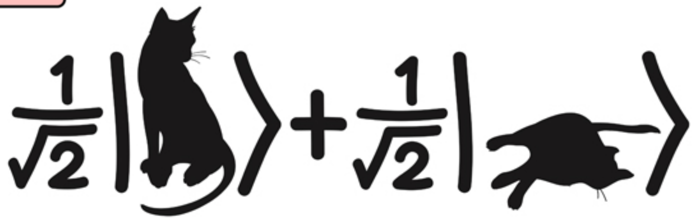

- Entanglement

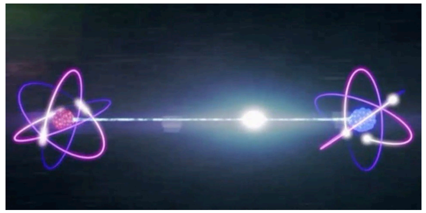

- Coherence

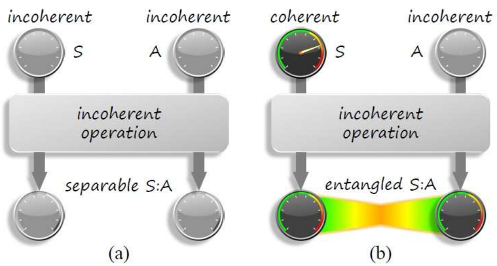

- Stochasticity

3. Speed

- Exponential search spaces of large problems appear a good match to quantum's inherent parallelism (superposition)
- Intelligence through fast learning
- HHL or Dwave as qGPU

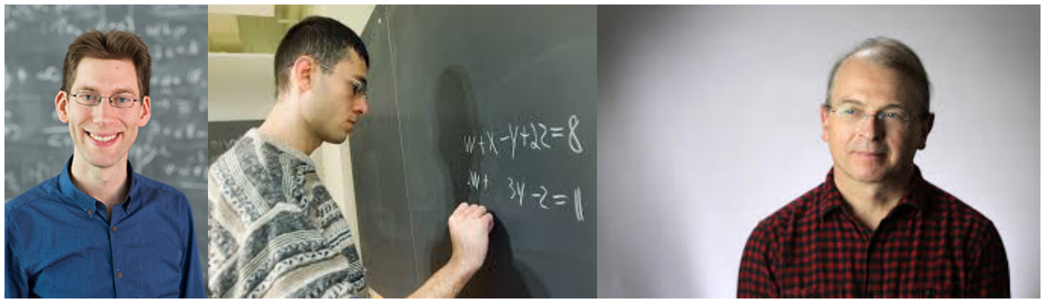

2. QML & classical data

- Quantum Recommender System for your books ar ads
- Quantum Image Recognition in your car
- Quantum Go
- Faster, better qML on classical data

Quadrant

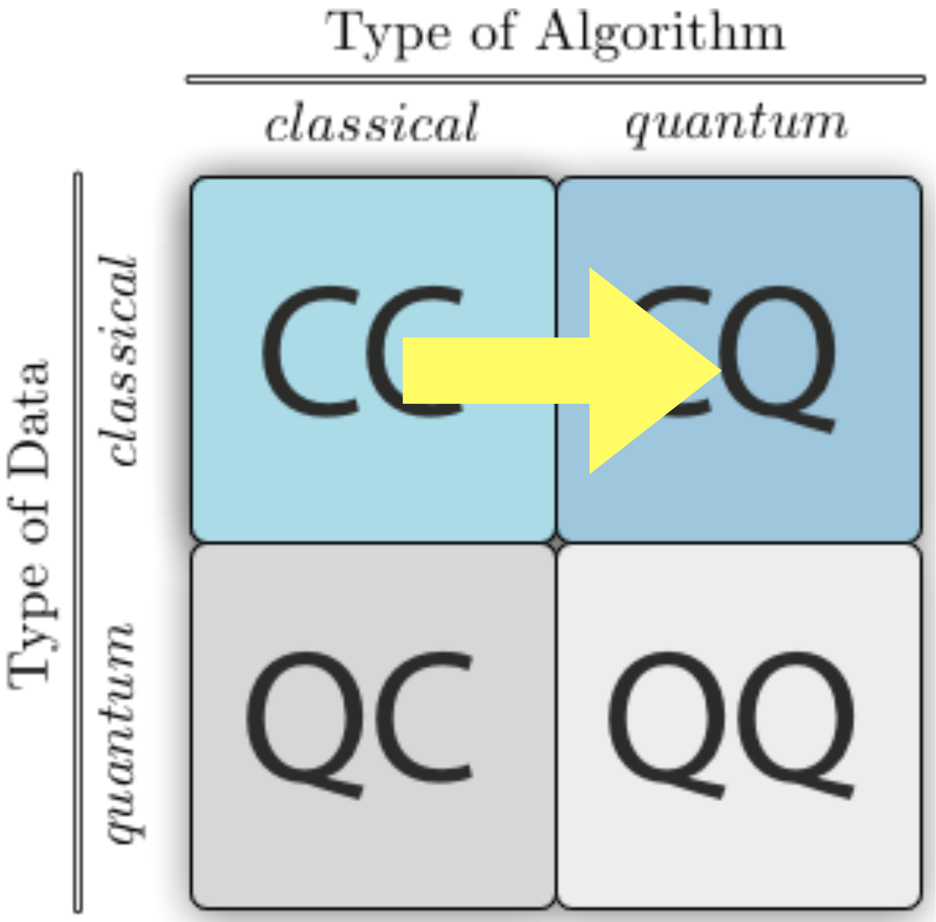

Type of Algorithm
classical
quantum
Type of Data
classical
quantum
CC
CQ
QC
QQ

Quantum ML Algorithms

- qSearch [Grover 1996]
- qPrincipal Component Analysis [Lloyd, Mohseni, Rebentrost 2013]
- qSupport Vector Machine [Rebentrost, Mohseni, Lloyd, 2013]
- qBayesian Network [Tucci 2012, Moreira, Wichert 2018]
- qHidden Markov Model [Clark, Huang, Barlow, Beige 2014, Cholewa, Gawron, Glomb, Kurzyk 2015, Srinivasan, Gordon, Boots 2017]
- qMatrix Operations, qBLAS [Harrow, Hassidim, Lloyd 2009, Le Gall 2012, Zhang, Zhang, Xue 2018, Shao 2018]
- qNeural Network [Schuld, Sinayskiy, Petruccione 2014, Neukart 2013, Ricks, Ventura 2004, Dorozhinsky, Pavlovsky 2018, Farhi, Neven 2018]
- qAnnealing [Behrman, Steck, Moustafa 2016, Li, Felice, Rohs, Lidar 2018]

Reviews (selection)

- [Schuld, Sinayskiy, Petruccione 2014]
- [Adcock, Allen, Day, Frick et al 2015]
- [Dunjko, Briegel 2018]
- [Arunachalam, De Wolf 2017]
- [Biamonte, Wittek, Pancotti, Rebentrost, Wiebe, Lloyd 2017]

2. The Data Problem

- Decoherence
- Supervised learning: get data in
- Transfer learning: get weights out

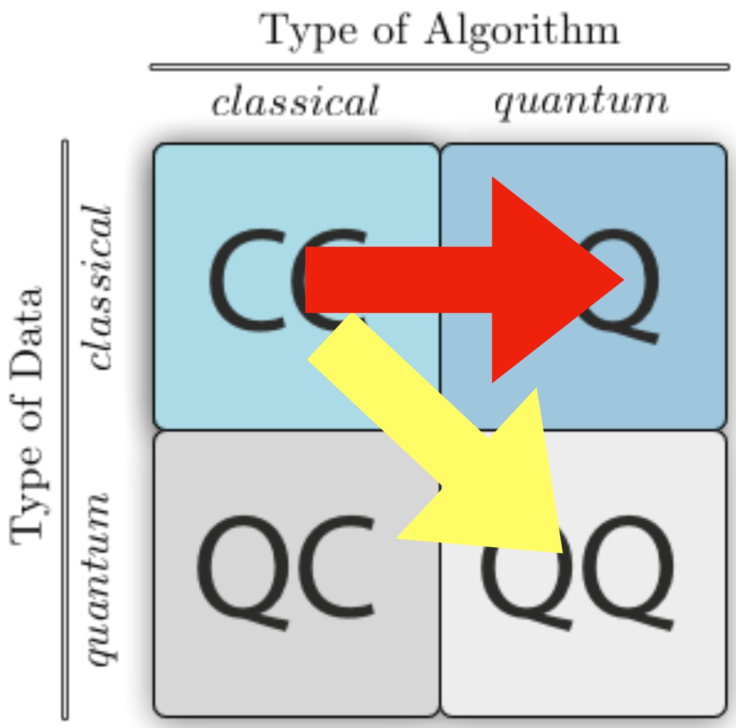

Observation

- Classical ML -> classical algorithms to achieve generalization

- QML: we are mimicing low level classical algorithms (NN, BLAS)
  Go for 3. Speed

- Runs into problems, such as decoherence, errors

- We could focus on high level goals: generalization
  use quantum effects as opportunity
  Go for 1. Algorithms

qFuture for AI?

- Transfer learning - decoherence problem
- Life long learning - decoherence problem
- One shot learning - superposition

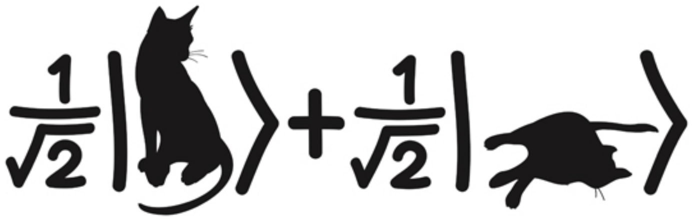

- Imperfect information - stochasticity

- Delayed credit assignment - entanglement
- Large problem spaces - superposition/annealing
- “quantum intelligence” - ???

Options for a QML agenda

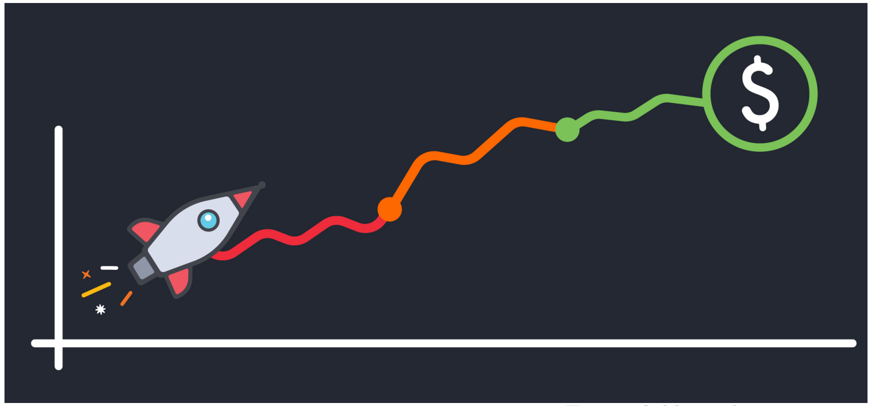

- speed up current AI
  (such as quantum deep learning)

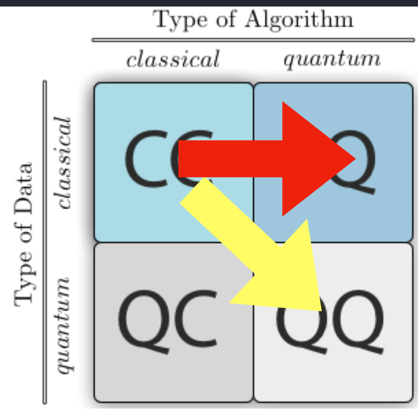

- speed up & work on future ML challenges
  (such as quantum meta learning)

- work on higher level AI approaches
  (such as entanglement and superposition for “new kinds of quantum intelligence: reasoning+generalization”)
  symbolic -> connectionist -> quantum

QML = fast NN?

QML agenda

- Go from 3. speeding up the past symbolic/connectionist learning

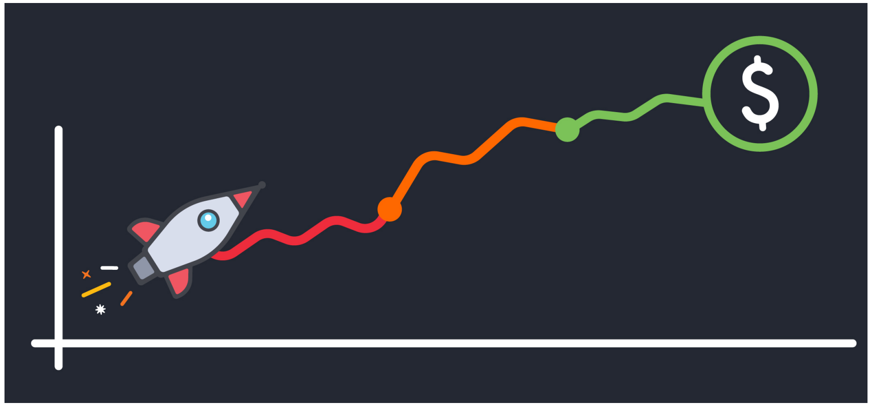

- To making new future quantum learning 1. Algorithms

Conclusion

- AI is hot. It has been cold
- Intelligence: Behaviorist approach. Intelligence is learning is generelization
- Breakthroughs in specific applications not in general approaches; Benchmarks work
- 1. Algorithms, 2. Data, 3. Speed
- Combination of two schools: symbolic and connectionist
- Copying classic AI in Q or new approaches to intelligence?
- QML: collaborate & understand quantum computing and AI communities

Universiteit Leiden

Limitations

- New developments in learning (imperfect information, transfer learning)
- Behavioral kind of intelligence. No emotion, no self awareness, no consciousness, no creativity
- AI. The semblance of intelligence by approximating large combinatorial spaces using machine learning

qFuture for AI?

- “QML” has focused on quantum versions of existing classical symbolic and connectionist ML algorithms AI as solving large combinatorial problems

    - (Data remains a problem)

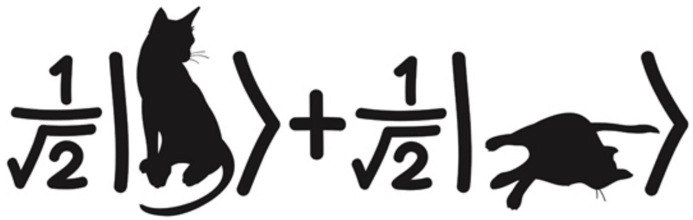

- QML could focus on list of Challenges in ML (sparse inputs, transfer learning, etc)

- Or QML could focus on its defining qualities (superposition, decoherence, entanglement) to “solve” learning/generalization