Artificial Intelligence DT8012

Introduction

Chapter 1, AIMA

What you'll learn from this course

An overview course – an introduction to AI technologies

- What is meant by Artificial Intelligence
  - what tools are used
  - what problems are targeted
- What problems can be solved with search
  - exactly or approximately
  - in dynamic or adversarial environment (game playing)
- How computers can represent knowledge
  - symbolic (logic) & non-symbolic (neural networks)
- How reasoning can be done with a machine
  - given knowledge that's certain & under uncertainty
- How a machine can learn & generalise

People

Slawomir Nowaczyk
Associate Professor
Course responsible

Taha Khan
Postdoc

Farzaneh Etminani
Postdoc

Hassan Nemati
PhD student

Labs and project

Stefan Byttner
Associate Professor

Reza Khoshangiz
Postdoc

Onur Dikmen
Senior Lecturer

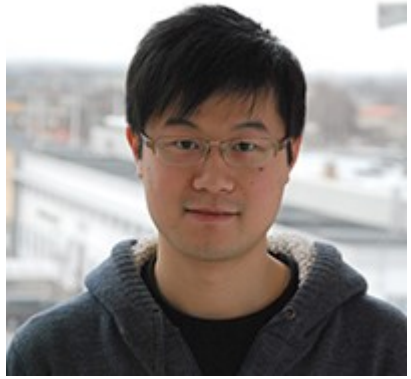

Yuantao Fan
PhD student

Labs

Course content follows the AIMA book closely

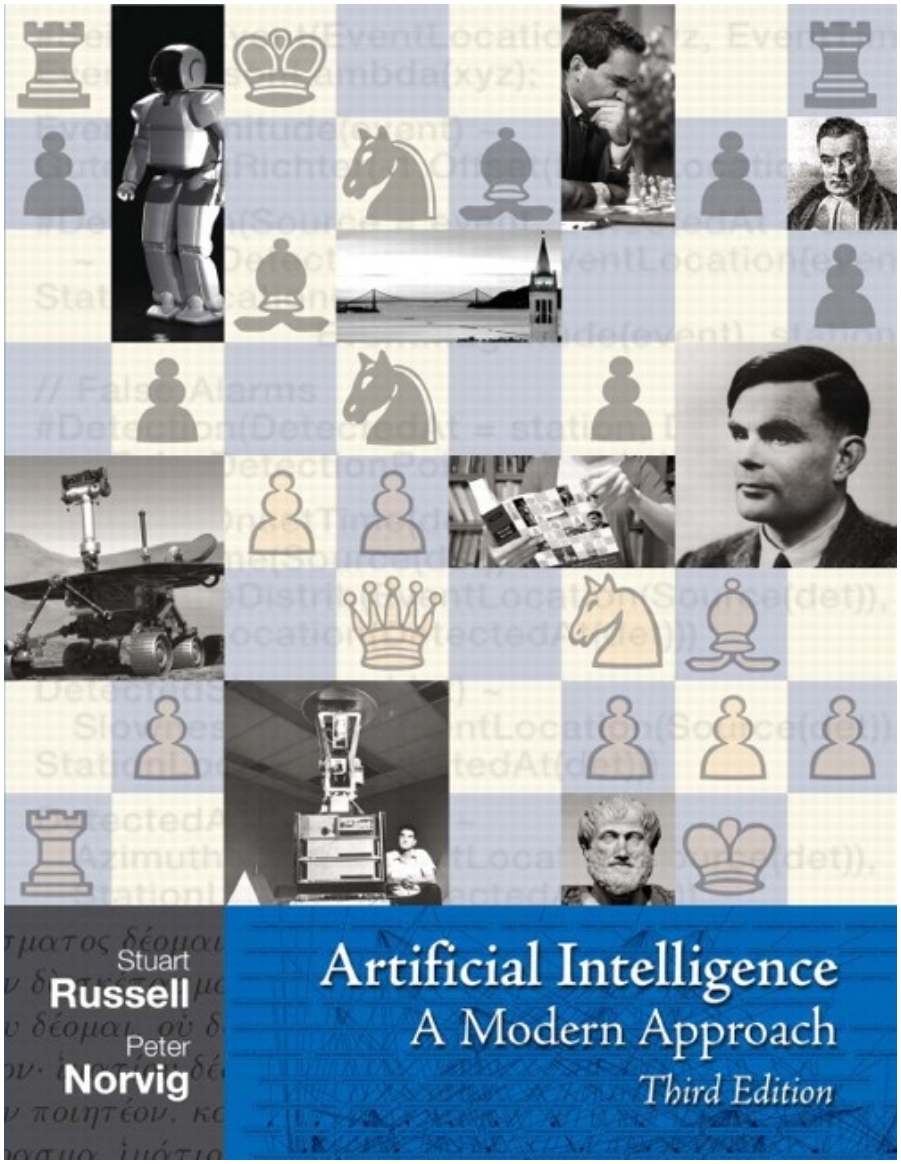

Time Management

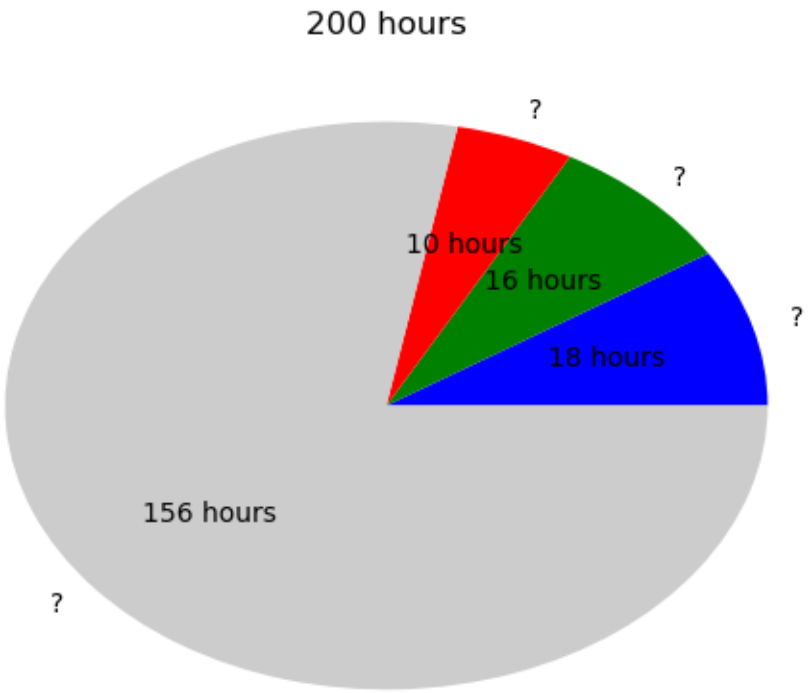

Time Management

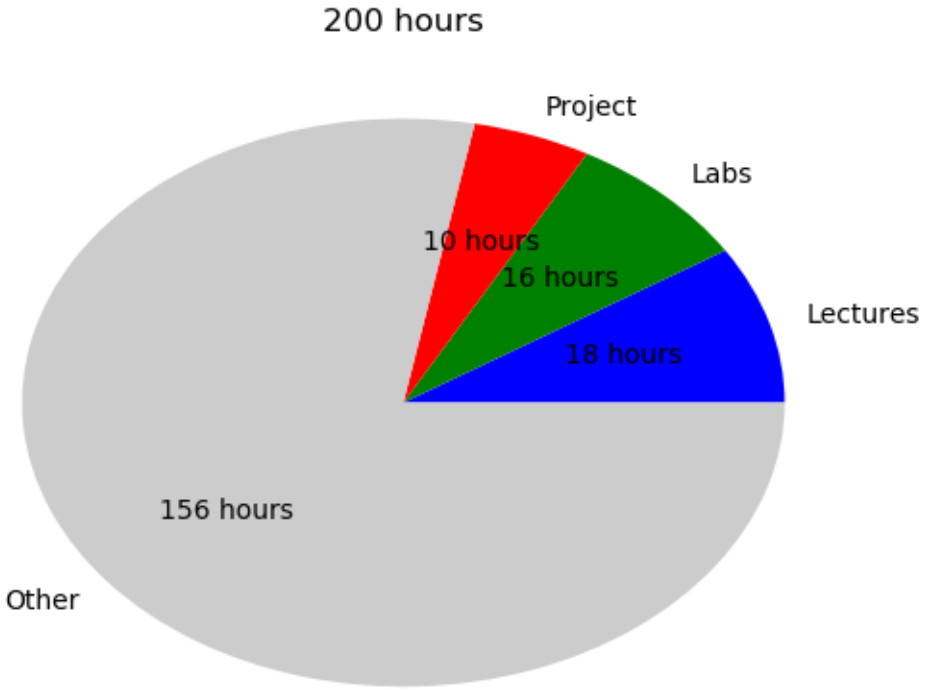

Course structure

7.5 credits course
10 weeks half-time
200 total work hours!

- Approximately 20 hours of lectures
- Lab work (4 labs, 4-8 hours each)
- Group project (40-80 hours)
  - implementation of poker-playing agents
  - tournament structure, two students per group
- Written + oral exam
- This is a difficult course, with a lot of material
  - it requires theory, practice & understanding

Course expectations

7.5 credits course
10 weeks half-time
200 total work hours!

- Topics are introduced during lectures & labs
  - AIMA book covers more topics than this course
- However not all details are given in the lectures
  - you need to read relevant book chapters in whole
  - and learn the complete material they cover
- The two exams have different focus
  - you need to both apply and explain what you know
  - memorisation is not enough, we aim for, and test during the exams, deep understanding of concepts

Course Plan

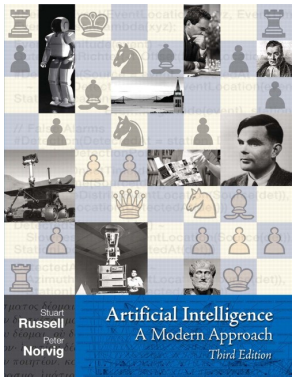

- Lecture 1: Introduction (chapters 1&2)
- Lecture 2: Uninformed Search (chapter 3)
- SEARCH
- Lecture 3: Adversarial Search (chapter 5)
- Lecture 4: Propositional Logic (chapter 7)
- LOGIC
- Lecture 5: First-Order Logic (chapters 8&9)
- Lecture 6: Uncertainty/Probability (ch. 13&16)
- PROBABILITY
- Lecture 7: Bayesian Networks (chapter 14)
- Lecture 8: Machine Learning 1 (chapter 18)
- LEARNING
- Lecture 9: Machine Learning 2 (chapter 18)

Course Material

- Course contents and relevant materials (e.g. lab descriptions) are available in Blackboard
- In particular, we will provide materials for you to prepare before each lecture (use it!)
- Lecture slides will be uploaded to Blackboard

Contact information

Stefan Byttner: stefan.byttner@hh.se Room E505
Slawomir Nowaczyk: slawomir.nowaczyk@hh.se Room E519
Yuantao Fan: yuantao.fan@hh.se Room E513
Hassan Nemati: hassan.nemati@hh.se Room E513

Questions?

Artificial Intelligence for (cooperating) Embedded Systems

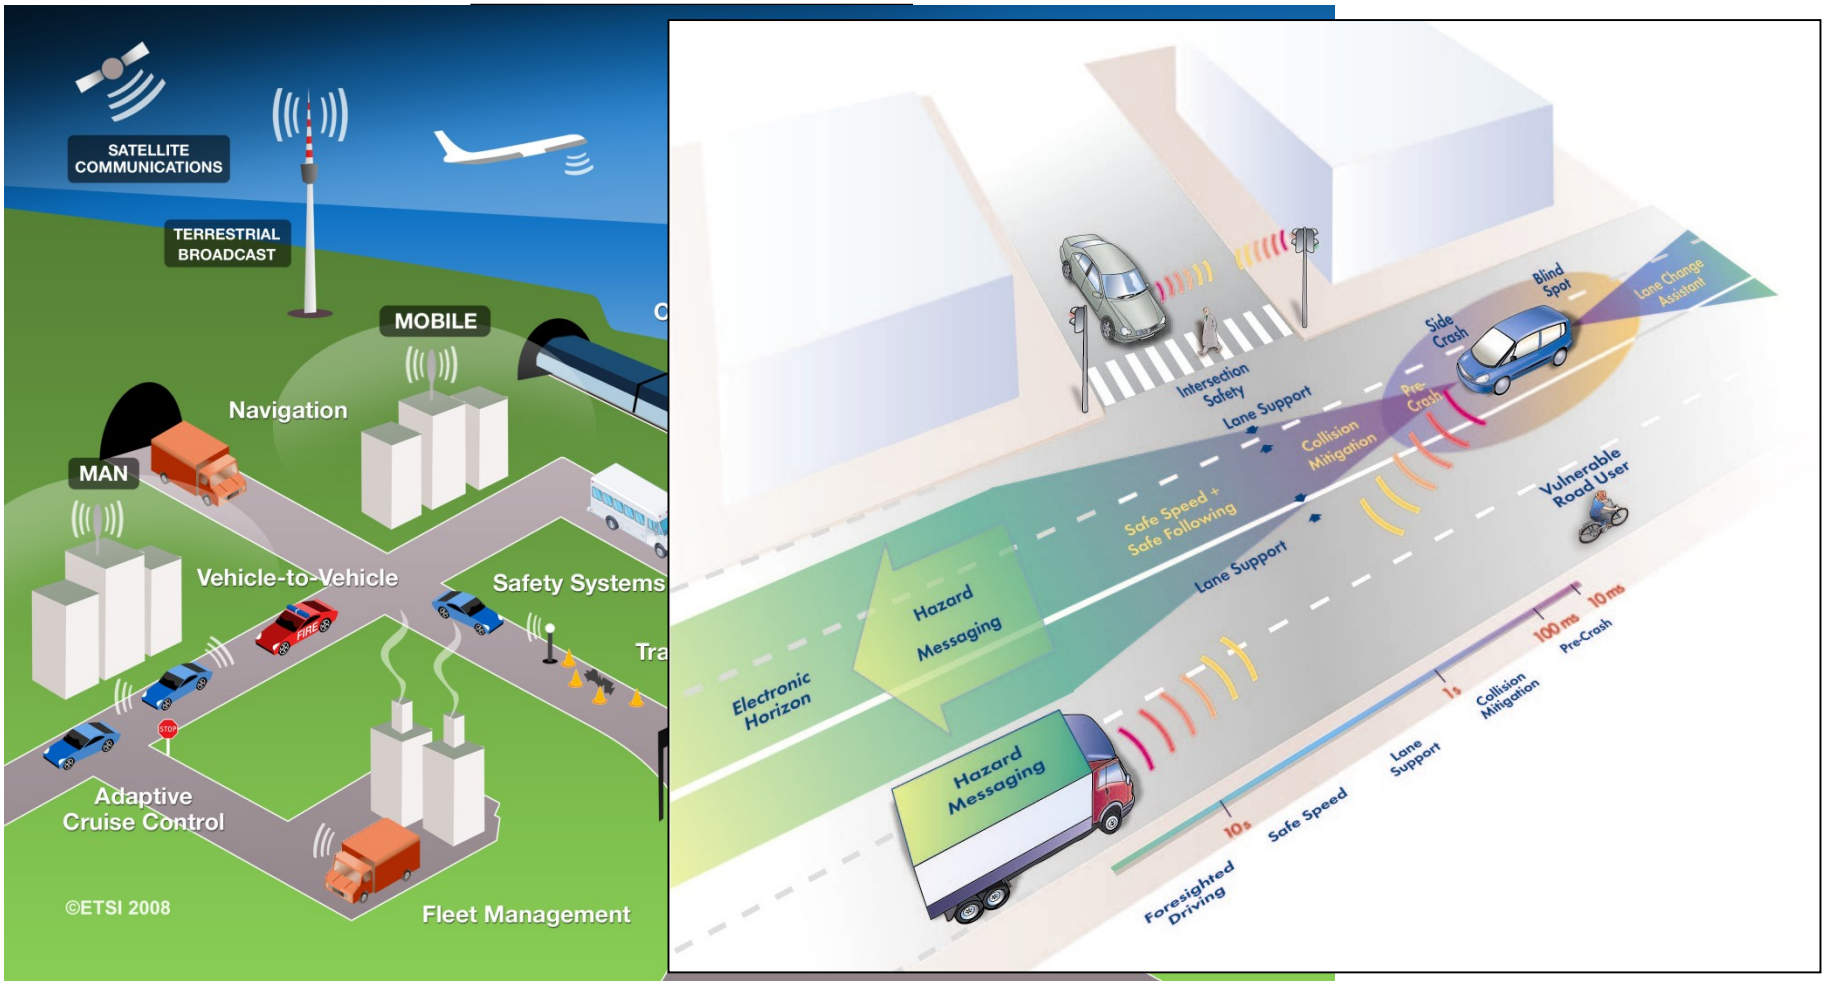

Intelligent embedded software and hardware for traffic control, safety, security, ...

Artificial Intelligence for (cooperating) Embedded Systems

University of Michigan and US Army
The Com-Bat: scavenge for power, stereoscopic cameras, microphones, detect radiation and airborne poisons.

Embedded intelligent systems for control of unmanned aerial vehicles

Artificial Intelligence for (cooperating) Embedded Systems

PARO – the robot seal. Keeps elderly company, like a pet.

de

a

8 other persons. It recognizes approximately 10,000 words and speaks spontaneously. Can perceive when something unusual occurs (alarm).

The ”pong” video arcade game

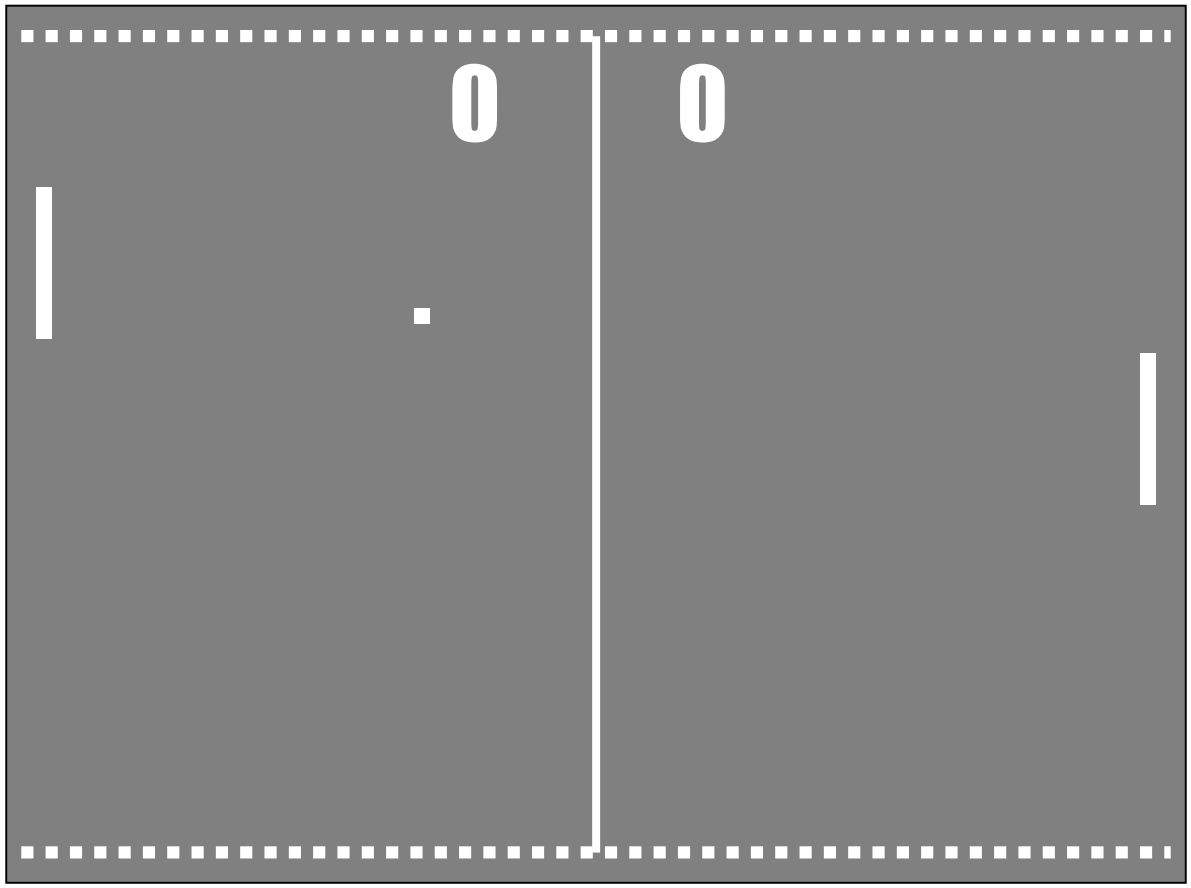

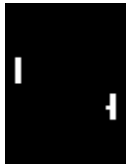

First public in 1972. The computer moves by calculating where the ball will cross the goal line and move the paddle there. Depending on difficulty, it sometimes does not move fast enough or moves to the wrong spot with some probability.

Games: Chess & IBM deep blue

- Deep Blue relies on computational power, search and evaluation.
- Deep Blue evaluates 200×10^6 positions per second. (Garry Kasparov evaluates 3 positions per second)
- The Deep Blue is a 32-node IBM RS/6000 SP with P2SC processors. Each node of the SP employs a single micro-channel card containing 8 dedicated VLSI chess processors, for a total of 256 processors working in tandem.
- Deep Blue calculates 100-200 billion (10^9) moves in three minutes.
- Deep blue typically searches 6 moves ahed but can go as far as 10-20 moves.
- Deep Blue beat the world champion Garry Kasparov in 1997 ”quantity has become quality”.

Deep Blue is “brute force”.
Humans (probably) play chess differently...

http://www.research.ibm.com/deepblue/meet/html/d.3.html

Games: Chess & IBM deep blue

- Deep Blue relies on computational power, search and evaluation.
- Deep Blue evaluates 200×10^6 positions per second. (Garry Kasparov evaluates 3 positions per second)
- The Deep Blue is a 32-node IBM RS/6000 SP with P2SC processors. Each node of the SP employs a single micro-channel card containing 8 dedicated VLSI chess processors, for a total of 256 processors working in tandem.
- Deep Blue calculates 100-200 billion (10^9) moves in three minutes.
- Deep blue typically searches 6 moves ahed but can go as far as 10-20 moves.
- Deep Blue beat the world champion Garry Kasparov in 1997 ”quantity has become quality”.

Deep Blue is “brute force”.
Humans (probably) play chess differently...

Note: in 1957, AI researchers thought that computers would beat the world chess champion within 10 years.

http://www.research.ibm.com/deepblue/meet/html/d.3.html

Do humans play chess differently?

Compare with HAL (the computer in ”2001: A Space Odyssey”). HAL plays ”tricky” and exploits the lower level of the opponent (the Astronaut Poole).

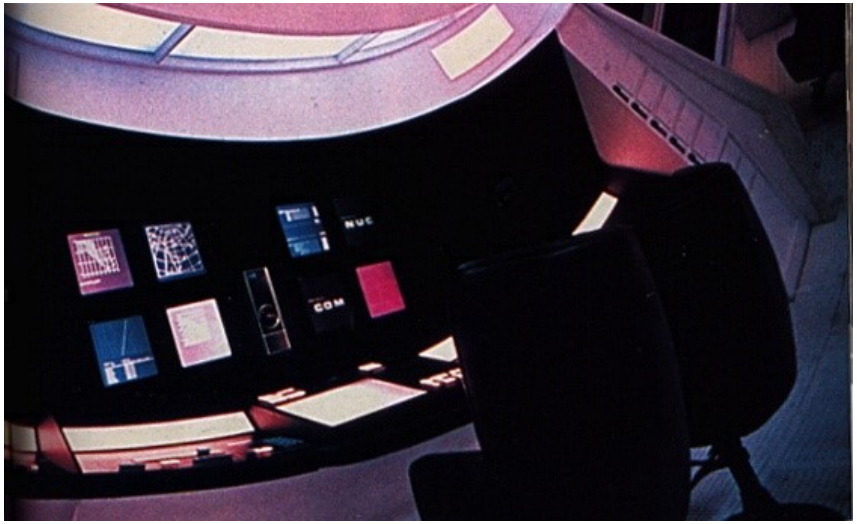

This is not ”computer-like”, but ”human-like”.

Computers, on the other hand, assume that the opponent will make the best possible move.

This is the minimax rule

Check out ”How HAL plays chess:
http://mitpress.mit.edu/e-books/Hal/chap5/five1.html

TD-Gammon

- The best backgammon programs use temporal difference (TD) algorithms to train a back-propagation neural network by self-play. The top programs are world-class in playing strength.
- 1998, the American Association of Artificial Intelligence meeting: NeuroGammon won 99 of 100 games against a human grand master (the current World Champion).
- TD-Gammon is based more on pattern recognition than search.

TD-Gammon is an example of machine learning. It plays itself and adapts its "rules" after each game depending on wins/losses.

http://satirist.org/learn-game/systems/gammon/

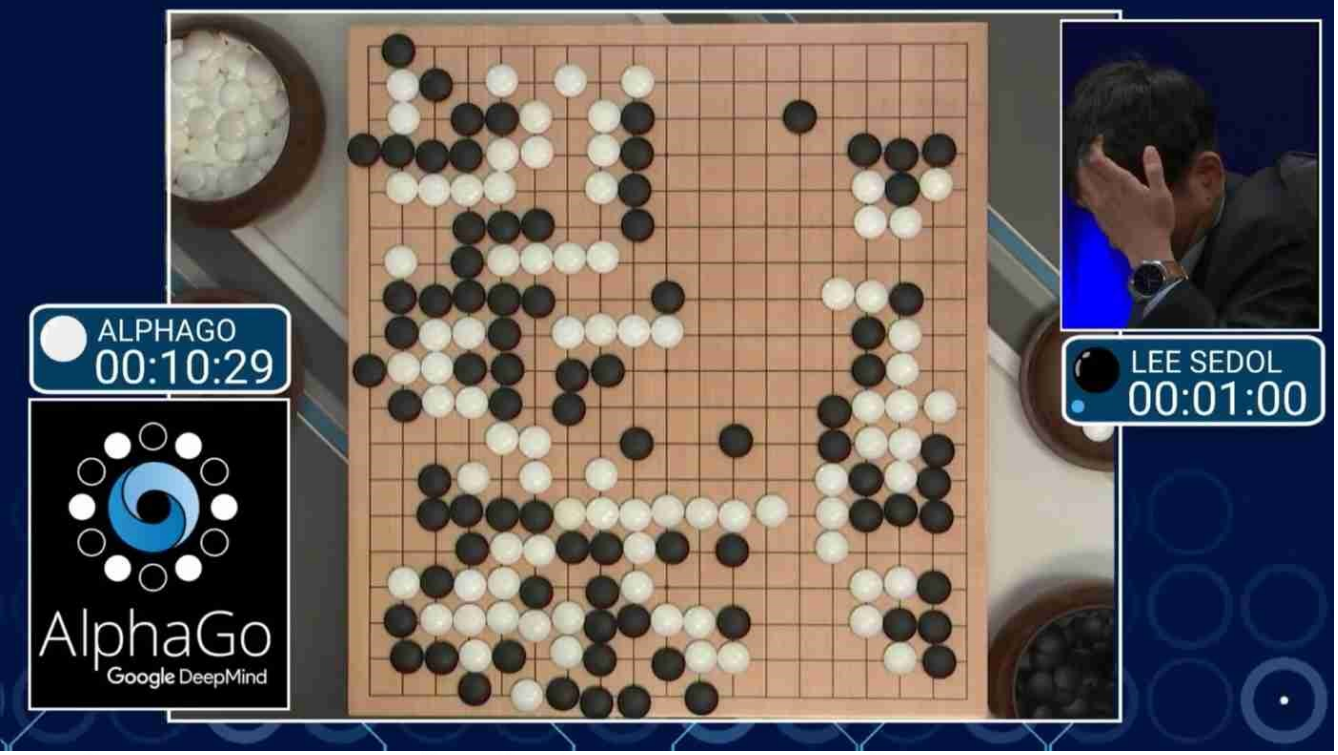

The story of AlphaGo so far

AlphaGo is the first computer program to defeat a professional human Go player, the first program to defeat a Go world champion, and arguably the strongest Go player in history.

AlphaGo's first formal match was against the reigning 3-times European Champion, Mr Fan Hui, in October 2015. Its 5-0 win was the first ever against a Go professional, and the results were published in full technical detail in the international journal, Nature. AlphaGo then went on to compete against legendary player Mr Lee Sedol, winner of 18 world titles and widely considered to be the greatest player of the past decade.

AlphaGo's 4-1 victory in Seoul, South Korea, in March 2016 was watched by over 200 million people worldwide. It was a landmark achievement that experts agreed was a decade ahead of its time, and earned AlphaGo a 9 dan professional ranking (the highest certification) - the first time a computer Go player had ever received the accolade.

> More on The Future of Go Summit in this video

AlphaZero is a computer program developed by research company DeepMind, which uses an approach similar to AlphaGo Zero's to master not just *Go*, but also *chess* and *shogi*. On December 5, 2017 DeepMind team introduced AlphaZero, which, within 24 hours, achieved a superhuman level of play in these games by defeating world-champion programs.

AlphaZero

From Wikipedia, the free encyclopedia

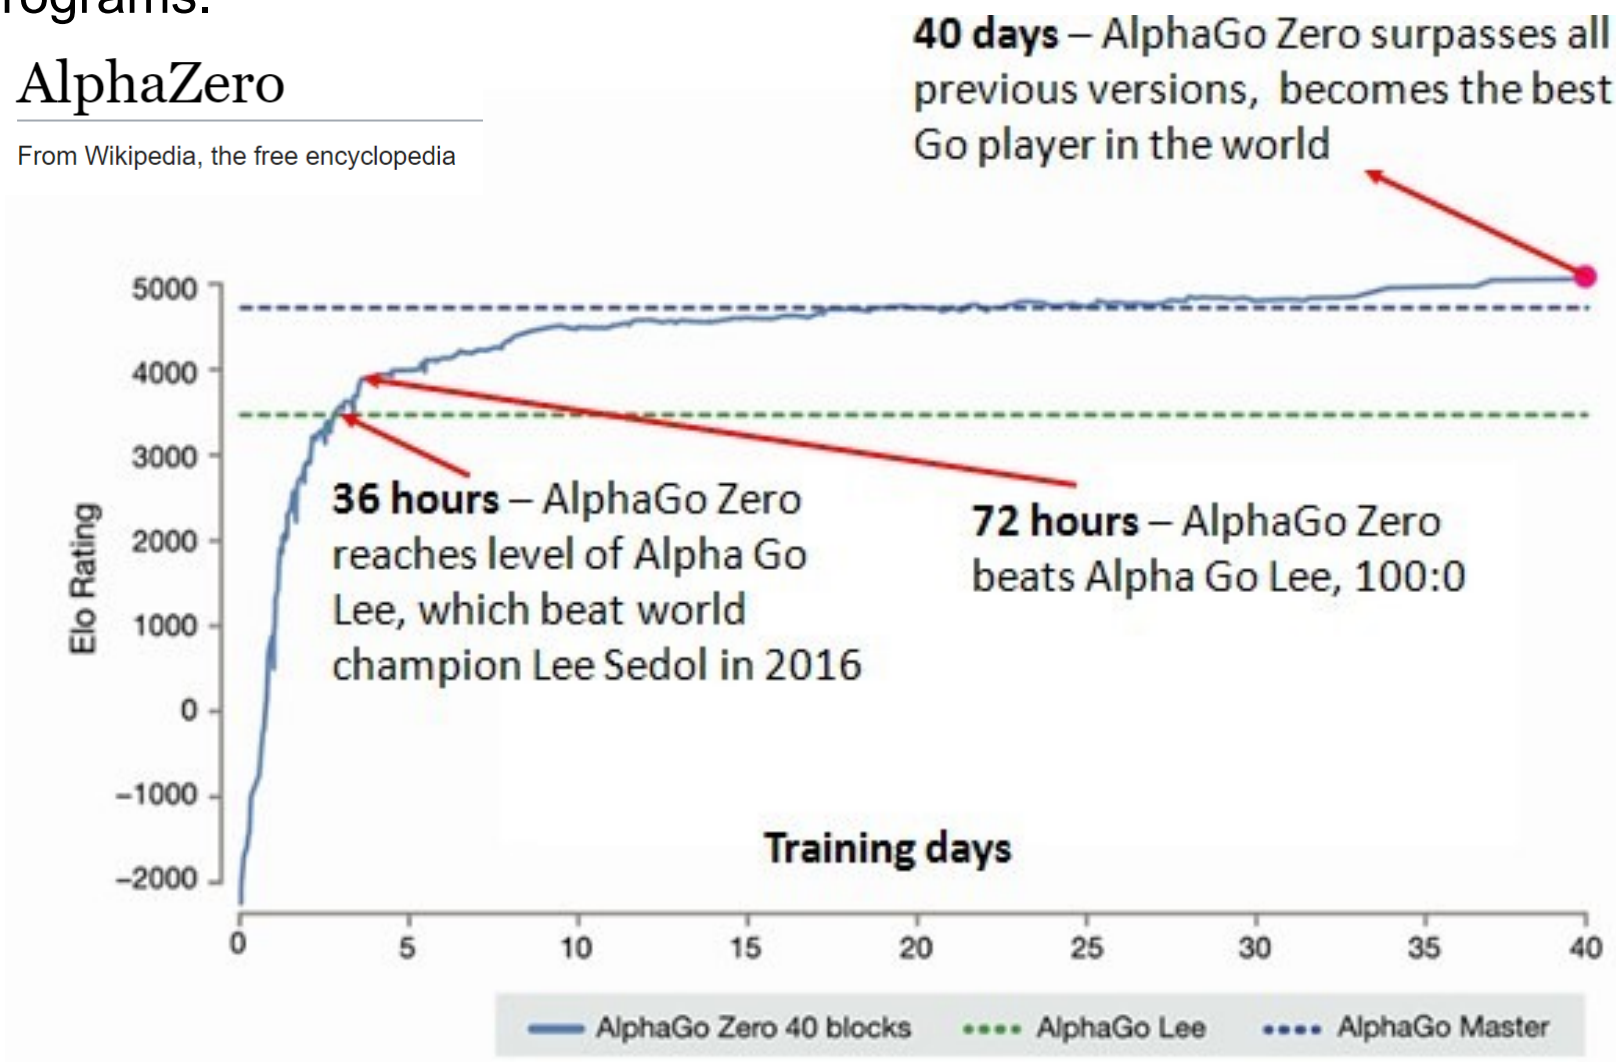

AI in video games

- ”Pong” was a first version…
- See online talk (Boston University) at http://www.bu.edu/buniverse/view/?v=1SaUoj65
- And at (UC Berkeley) http://www.youtube.com/watch?v=PsvsZuFgBzc
- Read tutorial (and watch slides) from Microsoft at http://research.microsoft.com/en-us/projects/ijcaiigames/
- Façade demo at http://www.youtube.com/watch?v=GmuLV9eMTkg

Search/query applications

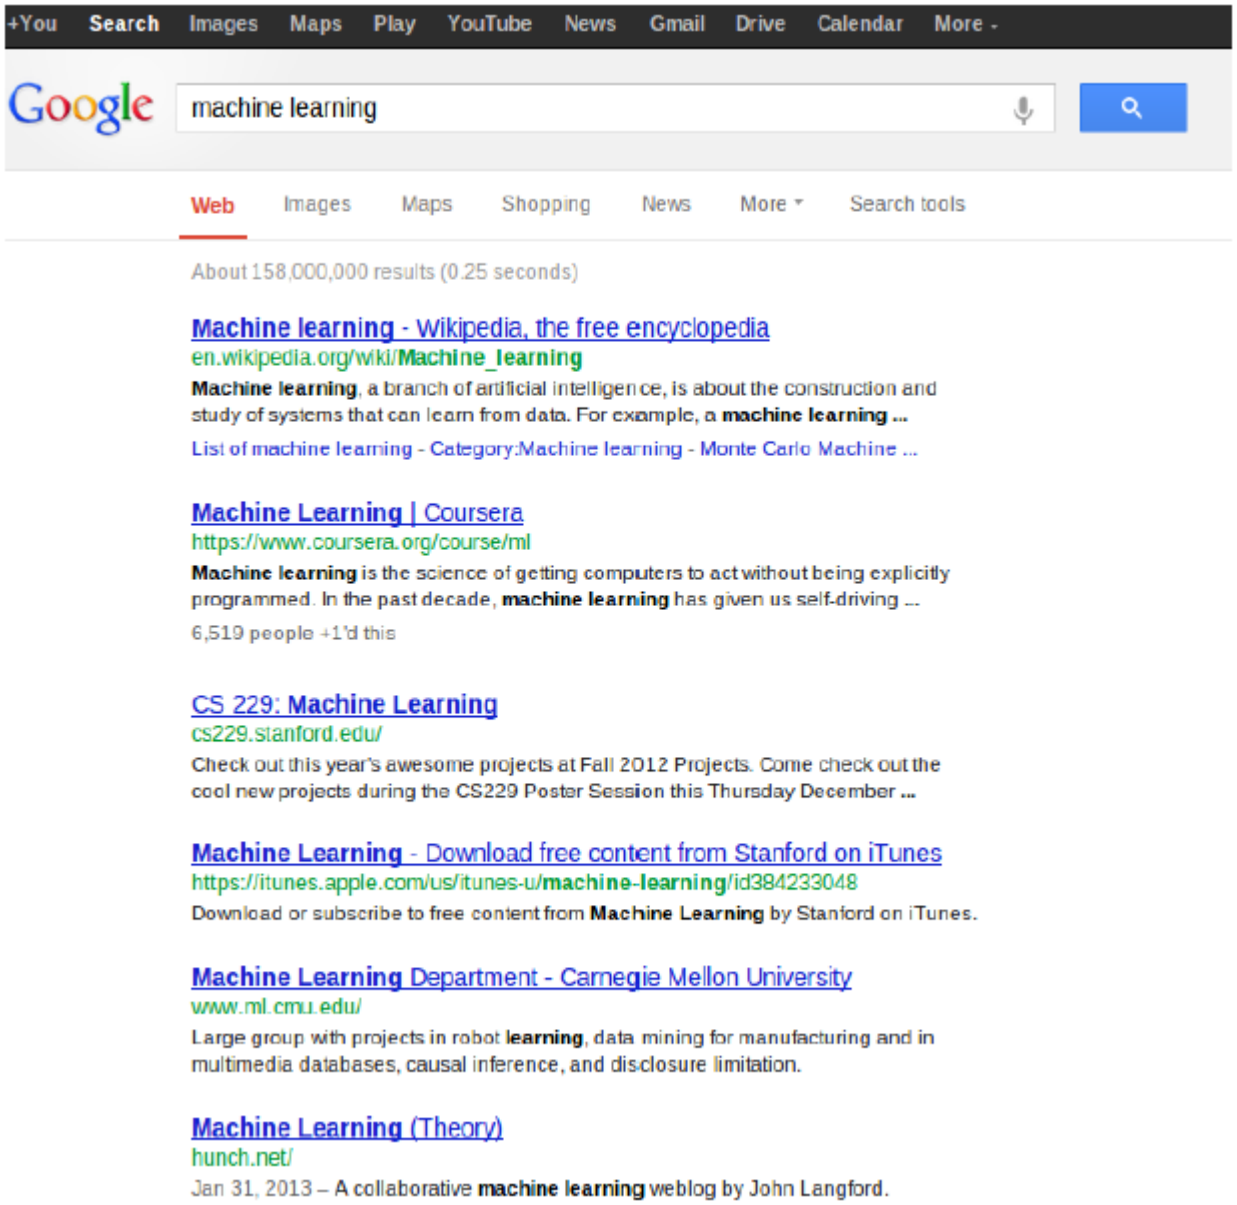

Machine translation

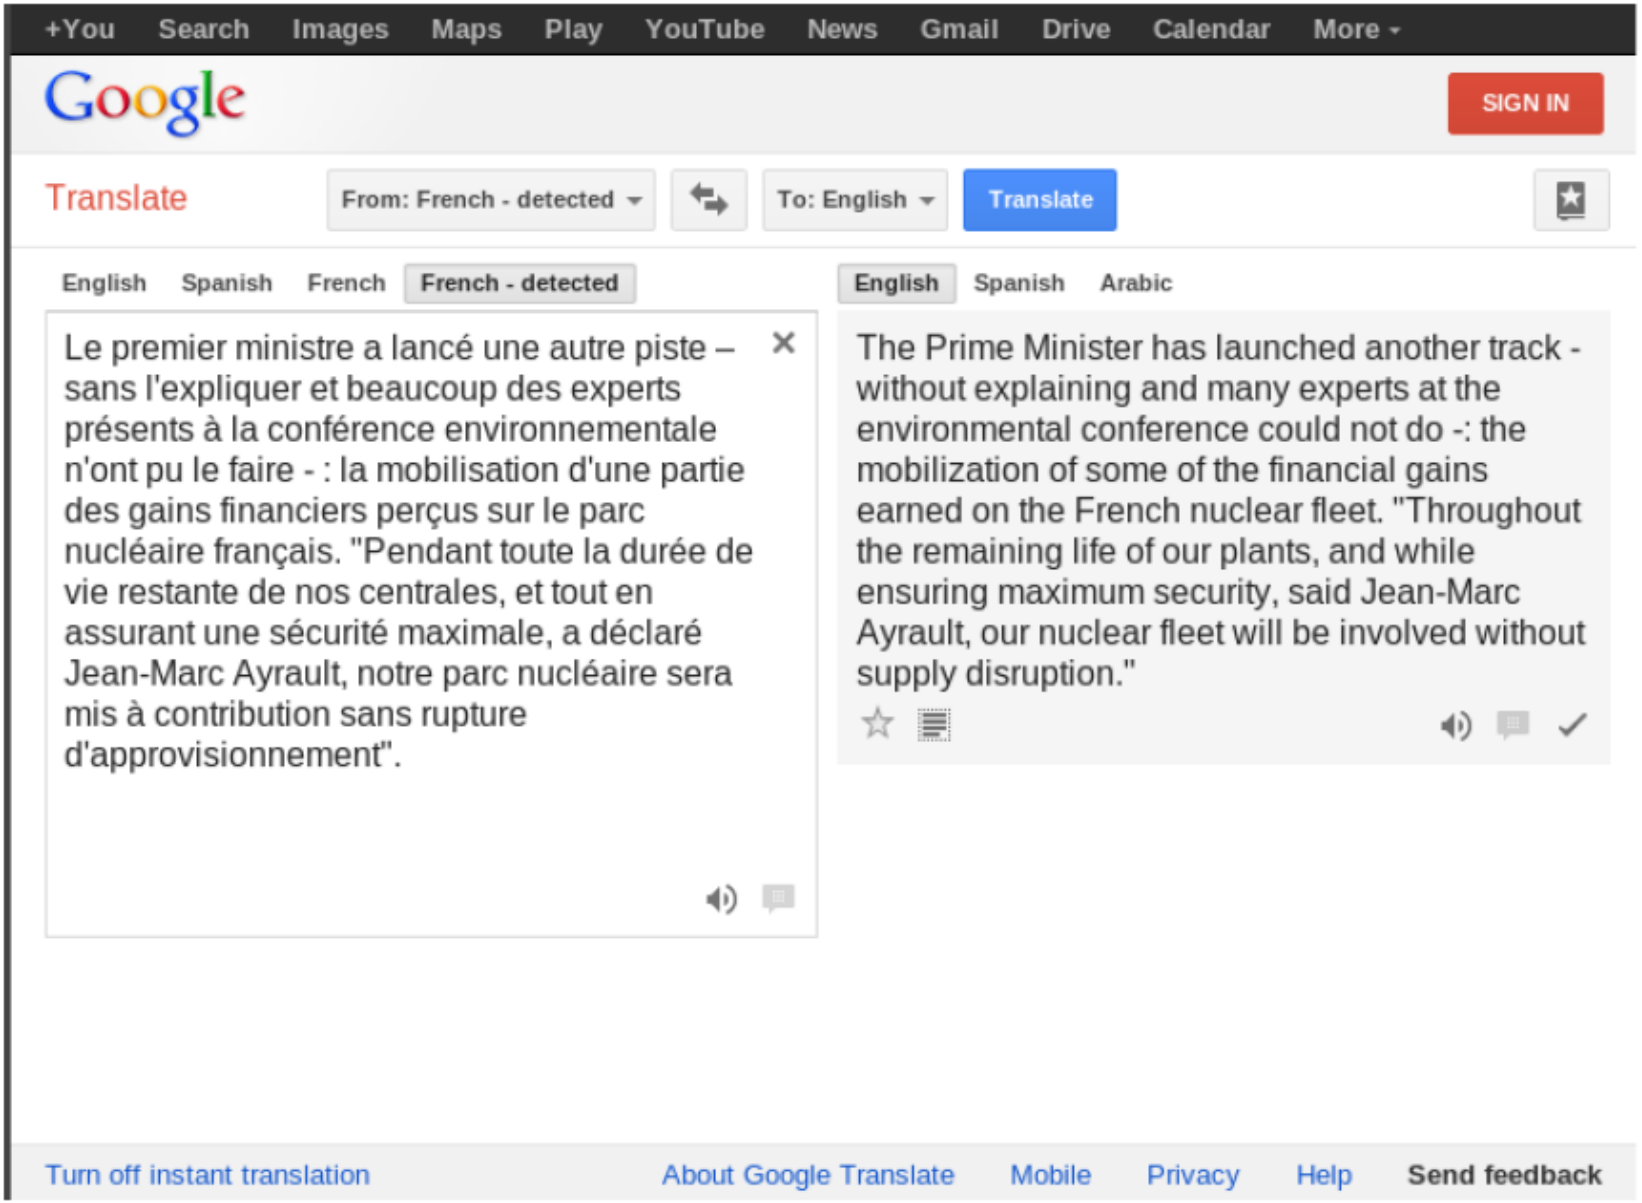

Route planning

Personal assistant

Navigating: Mars Autonomy Project

http://www.frc.ri.cmu.edu/projects/mars/dstar.html

Project at Carnegie Mellon, Pittsburgh

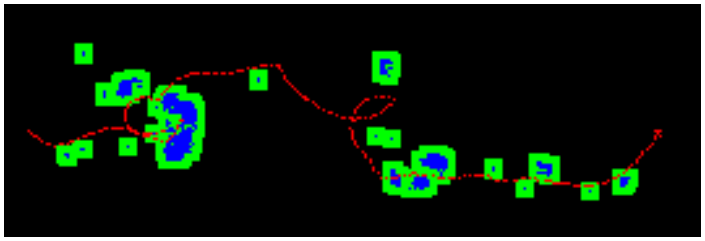

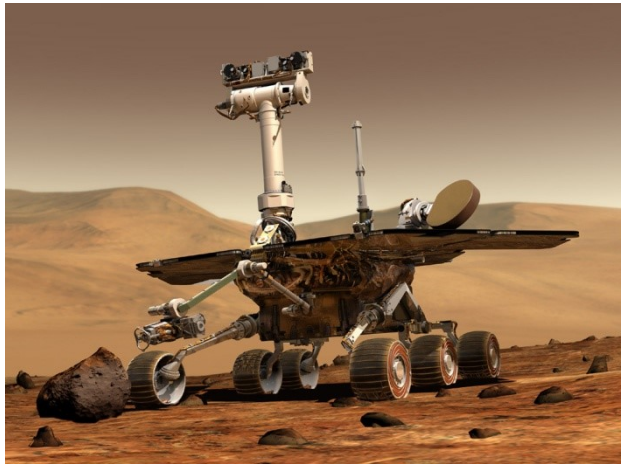

Project at JPL, Pasadena

Navigating: Under water

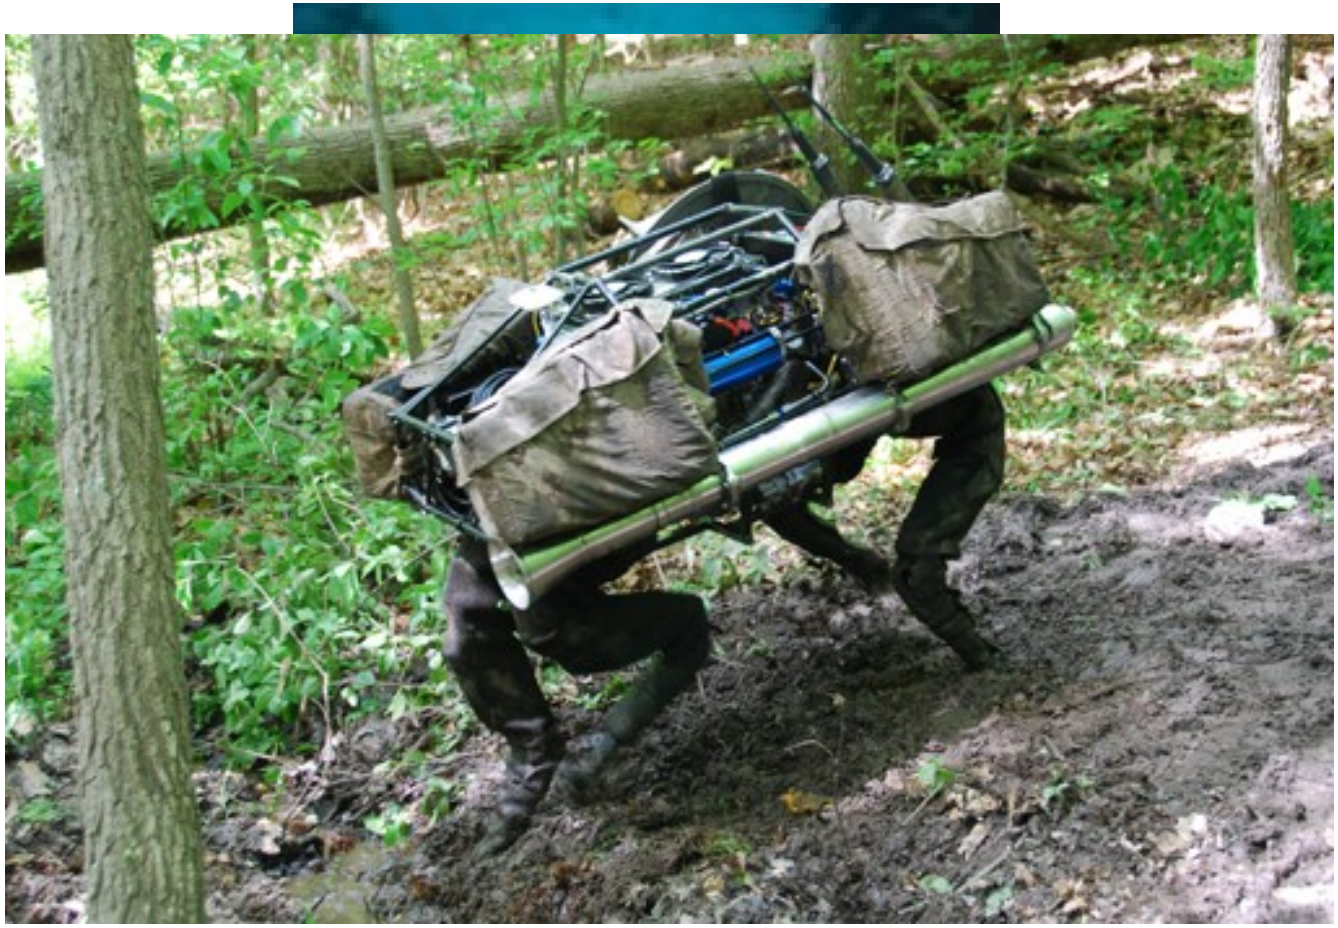

...and in the forest...

Autonomous driving

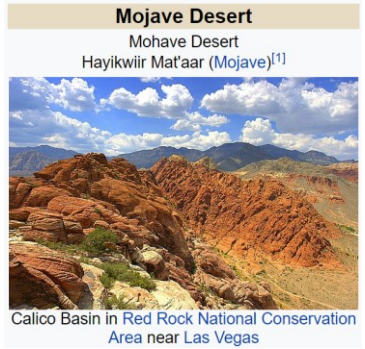

Stanley: The first car to finish the 2005 DARPA grand challenge.

Autonomous driving, 350 km in the desert.

It took 6 hrs and 54 min, with an average speed of about 50 km/h.

Stanford-group, lead by Prof. Sebastian Thrun.

Second was CMU, 10 minutes slower.

Nobody managed the similar task in 2004… the best car managed 11.78 km!

2007: Urban Challenge

- Autonomous driving 100 km in city environment in max 6 hours (about 15 km/h on average).
- Follow all traffic rules
- Manage other vehicles

2012: DARPA Robotics challenge

”… to develop ground robotic capabilities to execute complex tasks in dangerous, degraded, human-engineered environments.”

Unlike prior challenges, construction of the vehicles will not be part of the scope of the Robotics challenge. Boston dynamics have been contracted for developing 8 identical robots for the software teams. Robots are expected to be complete by Aug 2014.

Self-driving cars

$150,000 worth of equipment (including $70,000 LIDAR scanner)

Self-driving cars

The Drive Me trial starts in 2018 when 100 of our customers will drive XC90s equipped with the technology which will result in the IntelliSafe Autopilot, on Swedish roads

Self-driving cars

Grand Cooperative Driving Challenge (GCDC 2011)

11 teams from all over the world, Halmstad University took 2nd place

Grand Cooperative Driving Challenge (GCDC 2016)

10 teams from all over the world, this time Halmstad University won

AI fork-lift trucks (Halmstad)

Navigating: Vacuum cleaners

How do you guarantee that the vacuum cleaner doesn't get stuck and that it cleans the entire floor?

Small programs ~ 256 B

Navigating: helping elderly

And just helping out (housekeeping): http://www.youtube.com/watch?v=Uoq_r2dUf8g

Robots at home

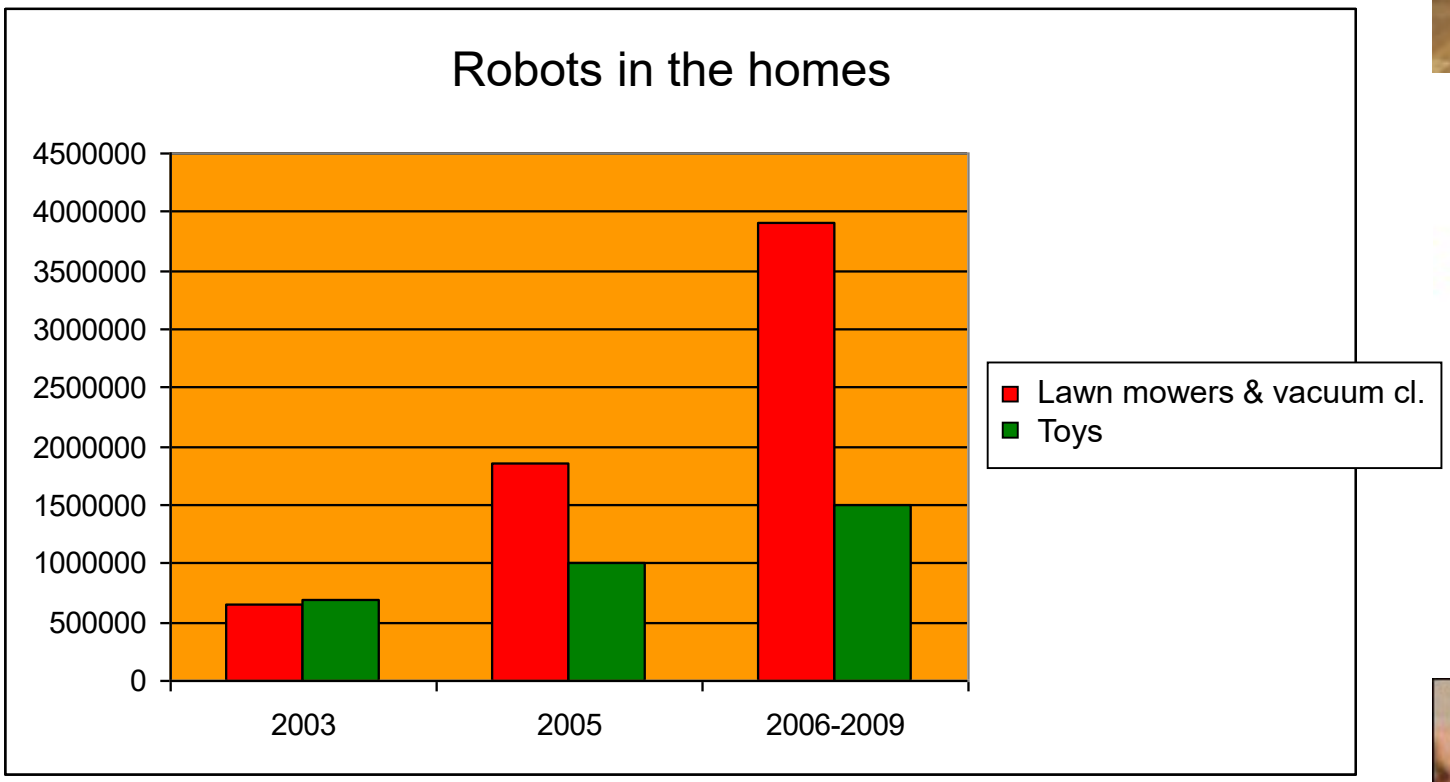

World Robotics Report 2004 & 2006

Robots at home

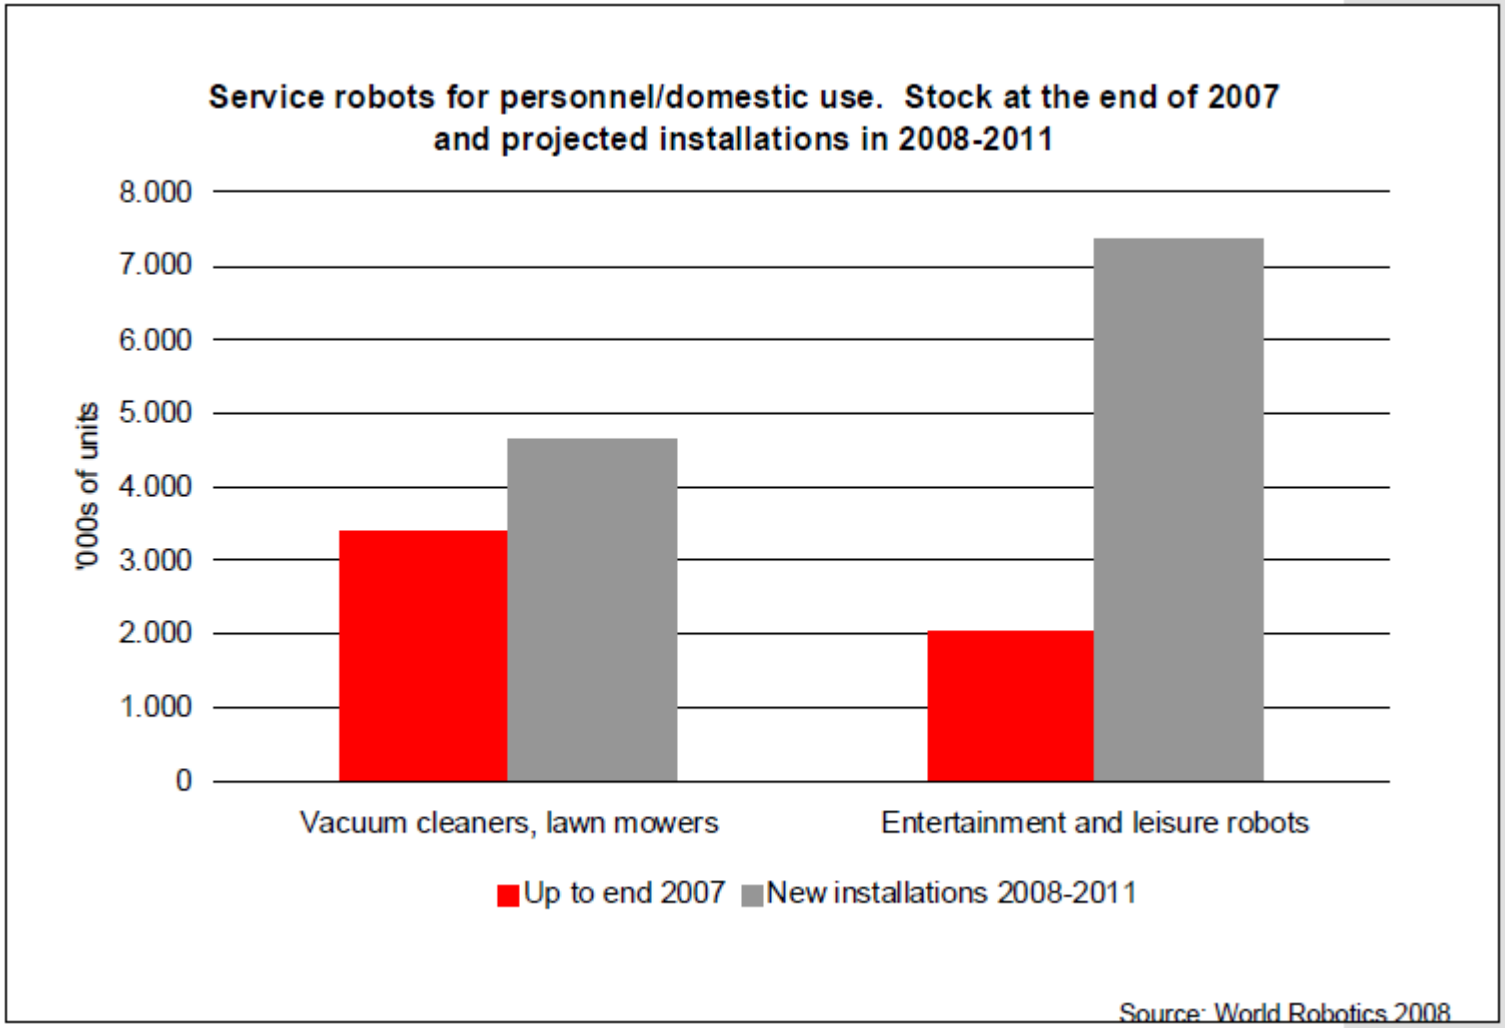

World Robotics Report 2008

Computers...

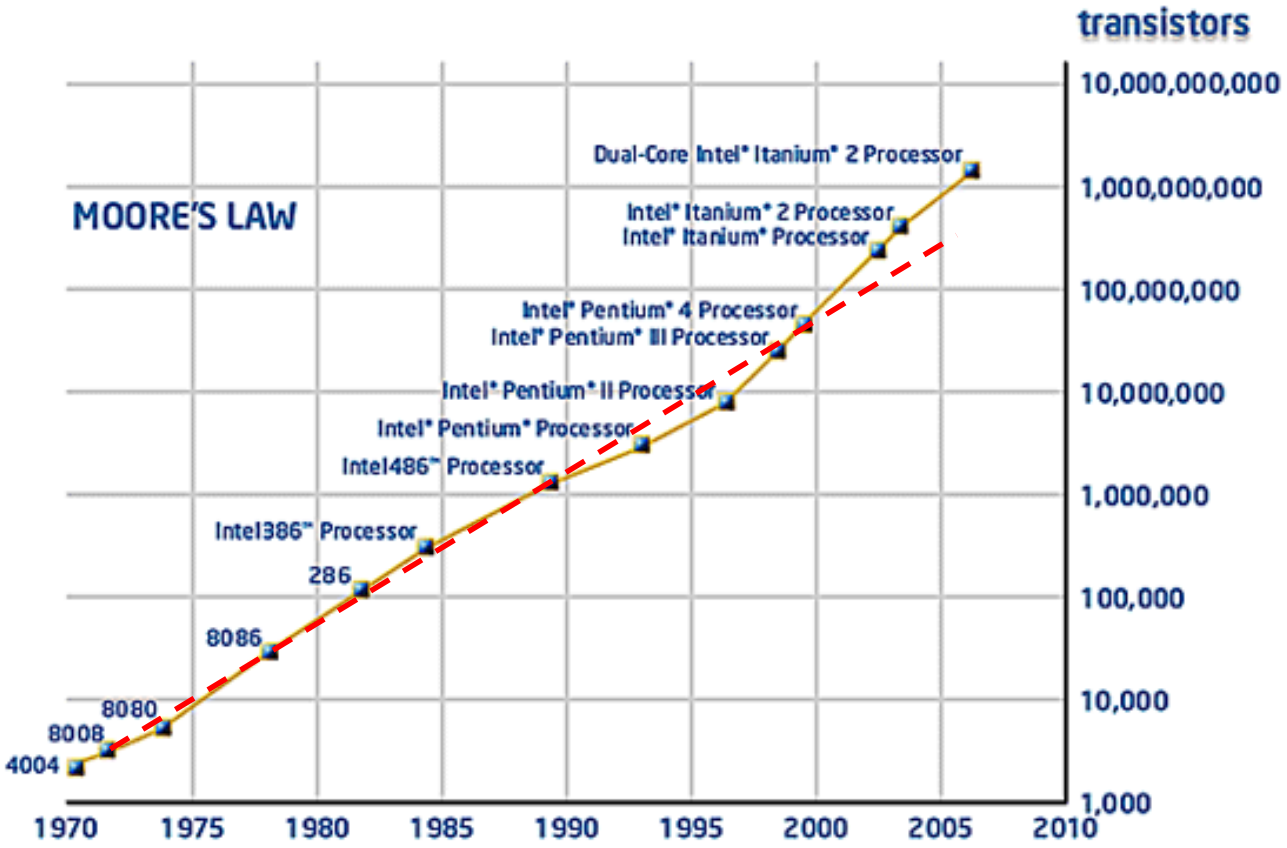
transistors
10,000,000,000
1,000,000,000
100,000,000
10,000,000
1,000,000
100,000
10,000
1,000
MOORE'S LAW
Dual-Core Intel® Itanium® 2 Processor
Intel® Itanium® 2 Processor
Intel® Itanium® Processor
Intel® Pentium® 4 Processor
Intel® Pentium® III Processor
Intel® Pentium® II Processor
Intel® Pentium® Processor
Intel486™ Processor
Intel386™ Processor
286
8086
8080
8008
4004
1970
1975
1980
1985
1990
1995
2000
2005
2010

...become cheaper and cheaper

A ”MIPS” becomes 1,000,000 cheaper in 40 years.

About half the price in one year.

(1 MIPS = 1 million ”instructions” per second)

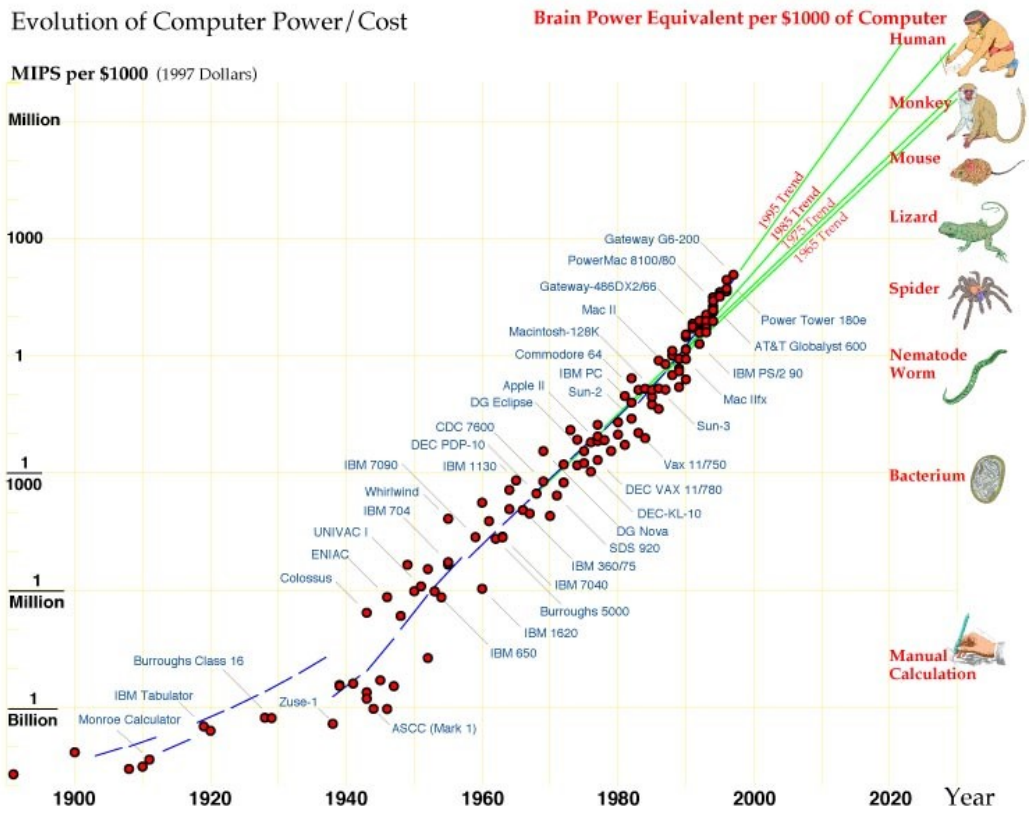

Image from Moravec

Computer memory becomes cheaper at a similar rate; half as expensive in two years.

What is AI?

- "The study of [rational] agents that exist in an environment and perceive and act." [Russell & Norvig]

What is AI?

Bellman, 1978 Thinking like a human	*Charniak & McDermott, 1985* Thinking rationally
Kurzweil, 1990 Acting like a human	*Luger, 2002* Acting rationally

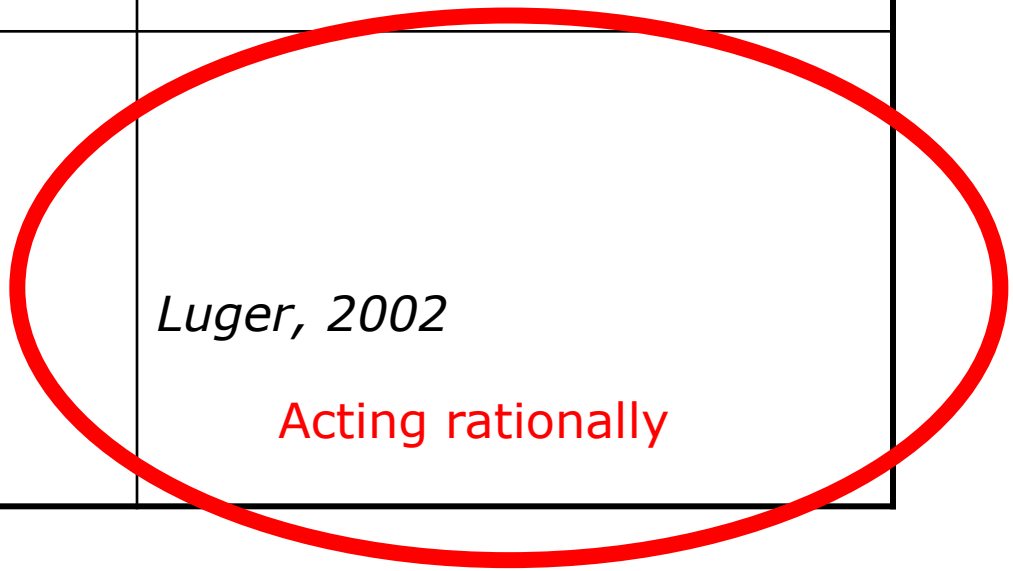

The Turing test – acting like a human

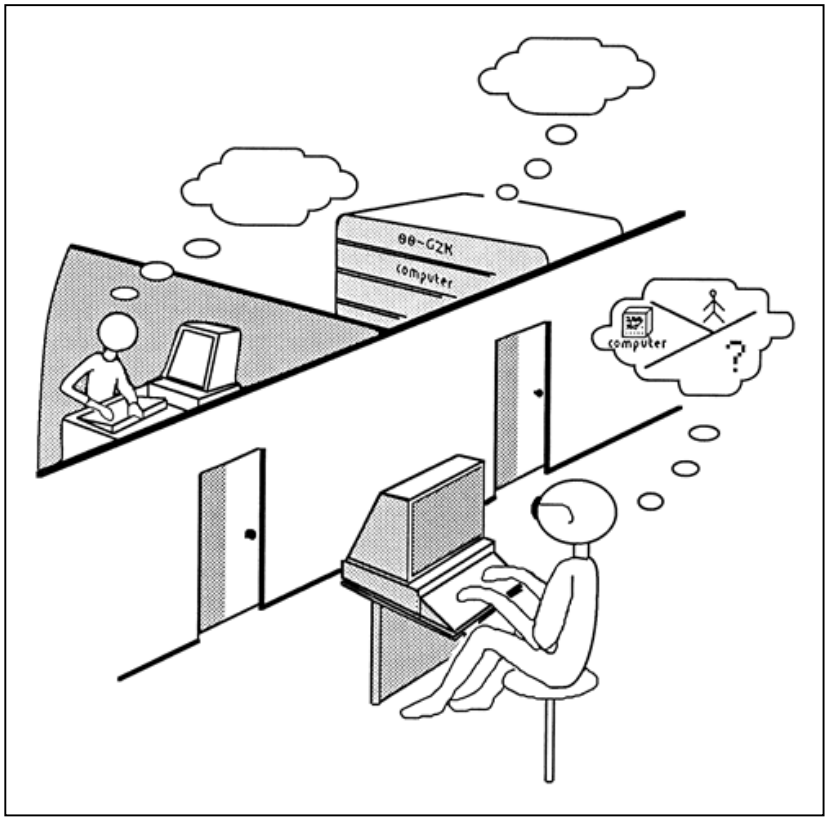

Suggested by Alan Turing in 1950.

If the interrogator cannot distinguish the human from the machine (robot), solely on the basis of their answers to questions, then the machine can be assumed intelligent.

Alan Turing

The Turing test provides...

- An objective notion of intelligence
  - no discussion on the ”true” nature of intelligence
- A way to avoid confusion by looking at how the computer reasons, or if it is conscious
- A way to avoid bias in favour of the human, by just focusing on the written answers

The Turing test can of course be generalized to other fields besides conversation

But it focuses too much on human behaviour. We are not trying to build humans (we already know how to do this...)

Problems with Turing test

- A test of the judge as well that of the actual AI machine
- Promotes imitators (con-artists)

See www.loebner.net

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

Chat bots:

http://www.abenteuermedien.de/jabberwock/index.php
http://www.alicebot.org/
http://www-ai.ijs.si/eliza-cgi-bin/eliza_script
http://www.simonlaven.com/

AI as ”rational agent”

- We will mainly focus on general principles of rational agents, in particular how can one approach constructing them
  - We define *rational* as ”achieving the best possible outcome”
  - It assumes the outcome is *measurable*
  - This definition is objective, and also general
  - We don’t have to meddle with what is ”human”

Fundamental issues in AI

- Sensing
  - How to extract relevant information from sensory input
- Representation
  - Facts about the world have to be represented in some way. Logic is one language that is used in AI. How should knowledge be structured? What is explicit, and what must be inferred? How to encode "rules" for inference so as to find information that is only implicitly known? How deal with incomplete, inconsistent, and probabilistic knowledge?
- Search
  - Many tasks can be viewed as searching a very large problem space for a solution. Use of heuristics and constraints.
- Inference
  - Some facts can be inferred from other facts.
- Learning
  - Learning is essential in an intelligent system.
- Planning
  - Starting with general facts about the world, about the effects of basic actions, about a particular situation, and a statement of a goal, generate a strategy for achieving the goal.