


Lecture 1B - What is AI?

Jesse Hoey
School of Computer Science
University of Waterloo

May 4, 2022

Readings: Poole & Mackworth 1.1

Understand principles and build devices that embody them



Can machines fly?



Can machines think?

Goals of Artificial Intelligence Research

- Scientific: Understand the Principles that make intelligent behaviour possible in natural or artificial systems
 - ▶ **analysis** of natural and artificial agents
 - ▶ **testing hypotheses** about the construction of intelligent agents
 - ▶ **building and experimenting with** intelligent agents
- Engineering: Design and build useful intelligent artifacts

History of AI: the Turing Test



Alan Turing: 1912-1954 was primarily known for codebreaking in WW11 (<https://www.imdb.com/title/tt2084970/>)

Also asked 'can machines think?', proposed the imitation game:

The interrogator is connected to one person and one machine via a terminal, therefore can't see her counterparts. Her task is to find out which of the two candidates is the machine, and which is the human only by asking them questions. If the machine can "fool" the interrogator, it is intelligent.

Conversations

ELIZA - (1966 Joseph Weizenbaum, Lisp)

ME: Hi

ELIZA: How do you do? What brings you to see me?

ME: I'm thinking of taking a course in artificial intelligence.

ELIZA: Why do you say that?

ME: I'm worried about it

ELIZA: Are you worried often?

ME: when I take AI courses

ELIZA: What happened when you take ai courses?

ME: I usually spend long nights doing assignments

ELIZA: Earlier you said you felt worried?

ME: about AI courses

ELIZA: Maybe your plans have something to do with this.

ME: no

ELIZA: Why not?

ME: because I think you are a computer

ELIZA: Is it because you think I am a computer that you came to me?

Loebner Prize - Turing test competition

- fool all judges: \$1m prize
- no one has ever won
- example interaction:

CONFEDERATE: Hello, my name is Sam and I am the human

JUDGE: Hi Sam, tell me which is larger, a TV screen or Belgium.

CONFEDERATE: Belgium is a little bit larger.

JUDGE: Wljso lwjejd wleje elwjeo wjeol, isn't it?

CONFEDERATE: Is that Polish?

JUDGE: Which is further away from this apartment, Queens or the Queen of England?

CONFEDERATE: If the Queen is visiting Manhattan, she is closer.

Loebner Prize

- Jabberwacky (Rollo Carpenter)
- just spits back nonsense,
- was very competitive for many years

JABBERWACKY: Hello, my name is George and I am the human.

JUDGE: If Isjkel lwehlx wouldn't you say?

JABBERWACKY: No, he's still good.

...

JUDGE: Which is further away from us, Queens or Inner Mongolia?

JABBERWACKY: HA! HA! That was stupid. Tell me a GOOD joke this time.

...

JUDGE: Which is larger, your brain or Belgium?

JABBERWACKY: Oh the one on the far left without a doubt.

Blocks World - SHRDLU - Winograd 1968

- first autonomous physical interaction
- blocks world only

Pick up a big red block

OK

grasp the pyramid

I DON'T UNDERSTAND WHICH
PYRAMID YOU MEAN

What does the box contain?


THE BLUE PYRAMID AND THE
BLUE BLOCK.

What is the pyramid supported
by?

THE BOX.

How many blocks are not in the
box?


FOUR OF THEM.




Game Playing

Two person, Zero-Sum games (competitive)


Game tree:




- 500 billion billion nodes (5×10^{20})
- Samuel 1950
- Jonathan Schaeffer, U. of Alberta (1992)
- <https://webdocs.cs.ualberta.ca/~chinook/play/>
- Heuristic search
- “solved” checkers - no search program could do any better
- World champion human-machine player



- Gerald Tesauro 1992
- IBM's Thomas J. Watson Research Center
- Neural Network
- Reinforcement Learning
- Just below human play






- game tree has more than 10^{100} nodes
- IBM - Deep Blue
- Heuristic Search - search depth: 7-8
- Deep Blue - Kasparov, 1996, Game 1
- Kasparov won 3-2-1
- Campbell, Hoane, Hsu, *Deep Blue* AI 2002
[https://doi.org/10.1016/S0004-3702\(01\)00129-1](https://doi.org/10.1016/S0004-3702(01)00129-1)
- Currently: Stockfish 10 <https://stockfishchess.org/>
- search depth of about 12
- AlphaZero: Silver *et al.* 2019
<https://doi.org/10.1126/science.aar6404>
beat Stockfish 10 155-6 (!)

- game tree has more than 10^{360} nodes
- Google Deep Mind : AlphaGo
- March 2016: AlphaGo beats Lee Sedol 4/5 games
- May 2017: AlphaGo beats Ke Jie 3/3 games
- <https://doi.org/10.1038/nature16961>



- Almost no domain knowledge
- Deep Reinforcement learning from pixels
- Convolutional Neural Networks
- better than human on 3/7 games
- arxiv.org/pdf/1312.5602v1.pdf



movie: <https://www.youtube.com/watch?v=V1eYniJ0Rnk>

- multi-agent problem
- imperfect information (partially observed map)
- large action space (10^8 possibilities)
- large state space
- delayed credit assignment
- Google Deep Mind :
<https://arxiv.org/pdf/1708.04782.pdf>



Article | Published: 30 October 2019

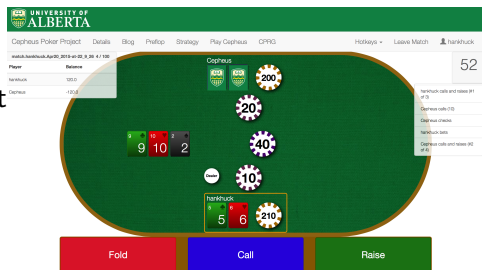
Grandmaster level in StarCraft II using multi-agent reinforcement learning

Oriol Vinyals , Igor Babuschkin, Wojciech M. Czarnecki, Michaël Mathieu, Andrew Dudzik, Junyoung Chung, [David H. Choi](#), Richard Powell, Timo Ewalds, Petko Georgiev, Junhyuk Oh, Dan Horgan, Manuel Kroiss, Ivo Danihelka, Aja Huang, Laurent Sifre, Trevor Cai, John P. Agapiou, Max Jaderberg, Alexander S. Vezhnevets, Rémi Leblond, Tobias Pohlen, Valentin Dalibard, David Budden, Yury Sulsky, James Molloy, Tom L. Paine, Caglar Gulcehre, Ziyu Wang, Tobias Pfaff, Yuhuai Wu, Roman Ring, Dani Yogatama, Dario Wünsch, Katrina McKinney, Oliver Smith, Tom Schaul, Timothy Lillicrap, Koray Kavukcuoglu, Demis Hassabis, Chris Apps & David Silver  - Show fewer authors

Nature **575**, 350–354(2019) | [Cite this article](#)

<https://doi.org/10.1038/s41586-019-1724-z>

- Michael Bowling *et al.*
- imperfect information
- Must model opponent
- Long-term payoff
- **Cepheus**
- CFR+: 4800 cores, 68 days: 900 core-years




Heads-up limit hold'em poker is solved Michael Bowling, Neil Burch, Michael Johanson, and Oskari Tammelin Science 9 January 2015: 347 (6218), 145-149.

<https://dx.doi.org/10.1126/science.1259433>

Video Game AI





Shakey
SRI
1970




José
UBC
2000



HRP-4C
AIST - Hiroshi Ishiguro
2010

Robotics



- Robocup 2017:
- <https://youtu.be/xkoXeF9oVH4>



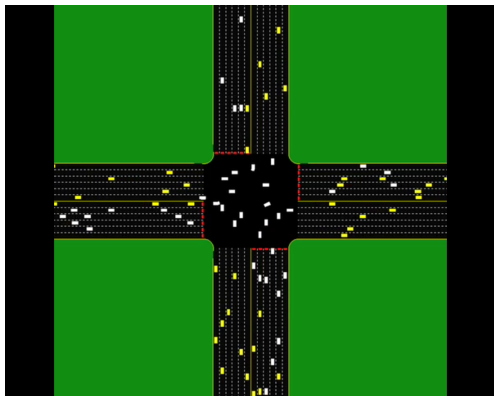
Autonomous Cars

- “Stanley” won the 2005 Darpa Grand Challenge
- Stanford/Sebastian Thrun



Autonomous Cars: 10 years?

- Peter Stone (UT Austin)
- <http://www.youtube.com/watch?v=4pbAI40dK0A>



More examples of AI in action

- space exploration
- disaster recovery
- web search
- advertising
- economy - predictions
- knowledge management, engineering
- circuit design, model checking, provability of systems
- air traffic control
- online selling and auctions
- social networks, computational social science and related disciplines
- ...



Are Self-Driving Cars (or AlphaGo/Deep Blue/etc...) Intelligent?

The synthesis and analysis of computational agents that act intelligently.

An agent acts *intelligently* when

- what it does is appropriate for its circumstances and its goals, taking into account the short-term and long-term consequences of its actions
- it is flexible to changing environments and changing goals
- it learns from experience
- it makes appropriate choices given its perceptual and computational limitations

Autonomous Cars: Flexible enough?



“They have to learn to be aggressive in the right amount, and the right amount depends on the culture.”

– Donald Norman, Design Lab, UCSD


from: New York Times “Google’s Driverless Cars Run Into Problem: Cars With Drivers”, 02/09/2015.

Autonomous Cars: Flexible enough?



theoatmeal.com/blog/google_self_driving_car

The Dawn of A.I. (1940s-50s)



Turing 1950



von Neumann 1944



Simon 1967

Phrenology 1880s




Limbic/Cortical Systems

Intermediate: paleopallium
Limbic System
Emotions



Rational Brain
Neocortex: neopallium
Higher thinking




Triune Brain



Primitive: archipallium
Survival, aggression

- Paul MacLean's *Triune Brain* 1960s
- limbic \approx hypothalamus, hippocampus, amygdala
- but these "systems" are really very mixed up in the brain

Antonio Damasio *Descartes' Error* Chapter 1



- Suffered brain damage (frontal lobe)
- Was perfectly good at reasoning and language
- Made disastrous decisions, or could not make decisions
- lacked “somatic markers” - “gut feelings” about decisions



1997: Rosalind Picard in *Affective Computing*
This book proposes that we give computers the ability to recognize, express and in some case “have” emotions. Is this not absurd?

Now:

- IEEE Transactions on Affective Computing
- International Conference on Affective Computing and Intelligent Interaction (ACII)
 - <https://acii-conf.net/2022/>
- Increasing awareness that emotions play a significant role in human intelligence
- but, still don't have “emotional machines” - why not?

Robots have feelings too..

- Darmstadt Dribblers:
- <https://www.youtube.com/watch?v=RbAlc-Y6j4o>





Emotions: the new AI

- Artificial Intelligence:
intelligence = **rationality**
- We now know that **emotions** are *necessary* for intelligence
- Emotions give “**heuristic**” social intelligence
- Encode a **social order** that allows us to work in a society



*With infinite resources,
are emotions necessary?*

The Singularity (von Neumann/Ulam)



Next:

- Agents (Poole & Mackworth chapter 1.3-1.10,2.1-2.3)
- Search (Poole & Mackworth chapter 3)