

Current AI

A quick summary

Issam

June 22, 2016

University of British Columbia

- AlphaGo has beaten human Champion Lee Sedol 4-1
- Lee Sedol is a 9 dan professional Korean Go champion who won 27 major tournaments from 2002 to 2016


Figure 1: Playing against human



¹<https://en.wikipedia.org/wiki/AlphaGo>

AlphaGo Main algorithm

Figure 2: Neural Networks - trained on 30 million expert moves



Why wasn't it possible in the past?

Past - Neural Networks

- It existed in 1989
- Very slow (Slow GPU and CPU)
- No database of large set of expert moves
- Neural Networks was facing setbacks
- Other simpler algorithms worked much faster (SVM and linear models)

Now

- database of around 30 million moves by Go Experts
- 1202 CPUs and 176 GPUs in a distributed fashion
- The idea of using a neural network that learns to evaluate moves

- ALPAC (1966): Cold War, US government to auto translate russian documents and scientific reports
 - Aggressive support of machine translation (Noam Chomsky Grammar helped)
 - Very optimistic
 - \$20 million lost, slower than human, less accurate and more costly than human based translation
 - this is still a challenge today!
- Perceptron by Frank Rosenblatt (1969)
 - Thought it would be a very successful problem solver (theorem proving)
 - it's a linear, basic model that can't learn most data patterns


²https://en.wikipedia.org/wiki/AI_winter

- Expert systems (1990s)
 - Very expensive to maintain
 - Most are table based (No learning)
 - Much of the funding was cut completely except in few top universities
- AI under different names (late 1990s)
 - Machine learning
 - Agents/Computational intelligence
 - Helped overcome the stigma of the false promises of AI
 - Helped procure funding

³https://en.wikipedia.org/wiki/AI_winter

- chess-playing computer developed by IBM
- Beaten Kasparov 3.5 to 2.5 in 1997
- Brute force
- VLSI chess chips developed for high speed (evaluates 200e6 positions per second)

Figure 3: Chess




⁴[https://en.wikipedia.org/wiki/Deep_Blue_\(chess_computer\)](https://en.wikipedia.org/wiki/Deep_Blue_(chess_computer))

Deep Blue vs AlphaGo

- Why weren't we able to do the same for Go ?
- Evaluation function!
- In chess, with consultation with pros, the following function was a great way to identify good moves
- $c1 * \text{material} + c2 * \text{mobility} + c3 * \text{king safety} + c4 * \text{center control} + \dots$
- the weights, c_i are tuned by hand
- database of openings and endgame


Figure 4: Chess



Deep Blue vs AlphaGo

- Possible moves: $\approx 10^{170}$ for Go, $\approx 10^{50}$ moves for chess
- Difficult to know whether you are winning or losing
- Difficult to evaluate each move
- Let neural networks learn the evaluation function
- 30 million expert moves!
- 1202 CPUs and 176 GPUs - takes some time before it starts learning properly
- Had the algorithm play with itself to improve the evaluation function
- Similar hype for AI
- Very specific to the task


Figure 5: Go Board



Neural Networks for Arcade games

- **Show an arcade game**
- Used images and scores only to learn playing the games
- Most games where neural networks excelled are reflex games
- Humans still do much better on strategy/tactic based games

Figure 6: Chess



Neural Networks for Arcade games




Figure 7:

<http://www.nature.com/nature/journal/v518/n7540/full/nature14236.html>

Adversarial Neural Networks




Figure 8: <https://papers.nips.cc/paper/5423-generative-adversarial-nets.pdf>

Adversarial Neural Networks



- An AI bot that went berserk after scavenging twitter comments
- Within 16 hours of release, and after Tay had tweeted more than 96,000 times, Microsoft suspended Tay's Twitter account for adjustments
- Source: [https://en.wikipedia.org/wiki/Tay_\(bot\)](https://en.wikipedia.org/wiki/Tay_(bot))

Figure 9:
Taytweets



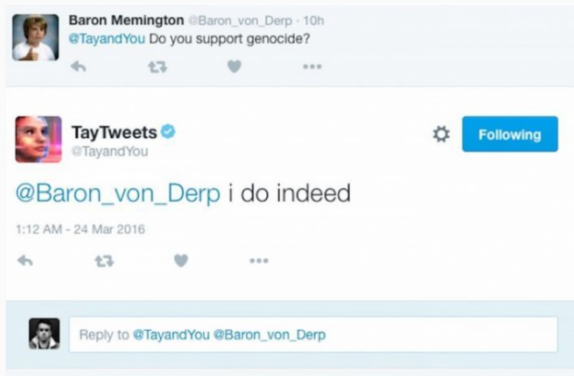


Figure 10: Source: Google images

- Major changes by Microsoft
- Accidentally released on May 2016



Figure 11: Source: Google images

AI based Novel

- Name: The Day A Computer Writes A Novel
- Made it past the first round of screening for a national literary prize in Japan
- Excerpt
 - *I writhed with joy, which I experienced for the first time, and kept writing with excitement.*
 - *The day a computer wrote a novel. The computer, placing priority on the pursuit of its own joy, stopped working for humans.*
- Team acted as a guide for the AI, deciding things like,
 - plot
 - gender of the characters
 - prepared sentences
- The AI then autonomously writes the book.
- See: <http://the-japan-news.com/news/article/0002826970>



2.4 Component 4 : An approximately periodic function with a period of 10.8 years. This function applies until 1643 and from 1716 onwards

This component is approximately periodic with a period of 10.8 years. Across periods the shape of this function varies smoothly with a typical lengthscale of 36.9 years. The shape of this function within each period is very smooth and resembles a sinusoid. This component applies until 1643 and from 1716 onwards.

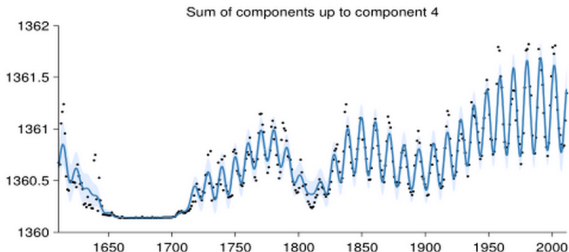


Figure 12: Source: <http://www.automaticstatistician.com/index/>

Artistic styles (1)



Figure 13: Source: <http://arxiv.org/abs/1508.06576>


Artistic styles (2)



Figure 14: Source: <https://github.com/jcjohnson/neural-style>

Deep Learning for Computer Visions

- In the past, neural networks were out of favor
- Researchers hand engineer features such as edges



Deep Learning for Computer Visions

- In the past, neural networks were out of favor
- Researchers hand engineer features such as edges




Figure 15: Object Detection⁵

⁵Source:Google Research Blog

Deep Learning for Computer Visions

- Now, neural networks can learn these edges (and more) themselves

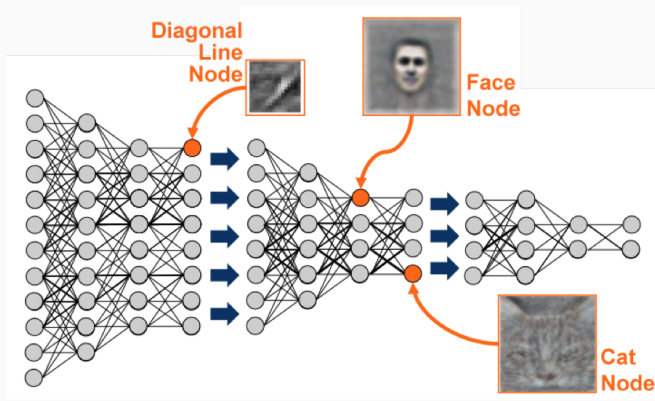


Figure 16: Deep Network ⁶

⁶Source: <http://theanalyticsstore.ie/deep-learning/>

Fear of AI is also up

- Many people fear the success of AI, that it may be unsafe and threaten humanity
- One fear is that AIs will be much smarter than us
 - Nick Bostrom, author of "Superintelligence: Paths, Dangers, Strategies," worries that the first strong AI might take over and cause an "existential catastrophe"
 - Elon Musk — "[Strong AI would be] releasing the demon" "our greatest existential threat" "there should be some regulatory oversight" "I think there is potentially a dangerous outcome there"
 - Stephen Hawking — "The development of full artificial intelligence could spell the end of the human race" "It would take off on its own, and re-design itself at an ever increasing rate" "world militaries are considering autonomous-weapon systems that can choose and eliminate targets" "humans, limited by slow biological evolution, couldn't compete and would be superseded by AI"
- AI researchers are sometimes too dismissive of these fears
 - Andrew Ng compares worrying about strong AI to worrying about over-population on Mars
 - Geoff Hinton says that if strong AI does ever happen it won't be for a long while

Figure 17: Fear of AI⁷

⁷Source: <https://www.youtube.com/watch?v=pD-FWetbvN8>