

Three waves of AI

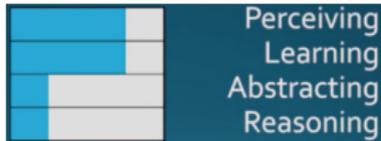
DARPA in 2017 identified the following three waves of AI:

- *First wave: Handcrafted knowledge.* Essentially the symbolic paradigm.
- *Second wave: Statistical learning.* Essentially the subsymbolic paradigm.
- *Third wave: Contextual adaptation.* Essentially the combination of symbolic and subsymbolic methods, allowing **learning**, **logical reasoning** and **explainability**.

The third wave is receiving increased attention and research focus.



First wave:
Handcrafted knowledge



Second wave:
Statistical learning

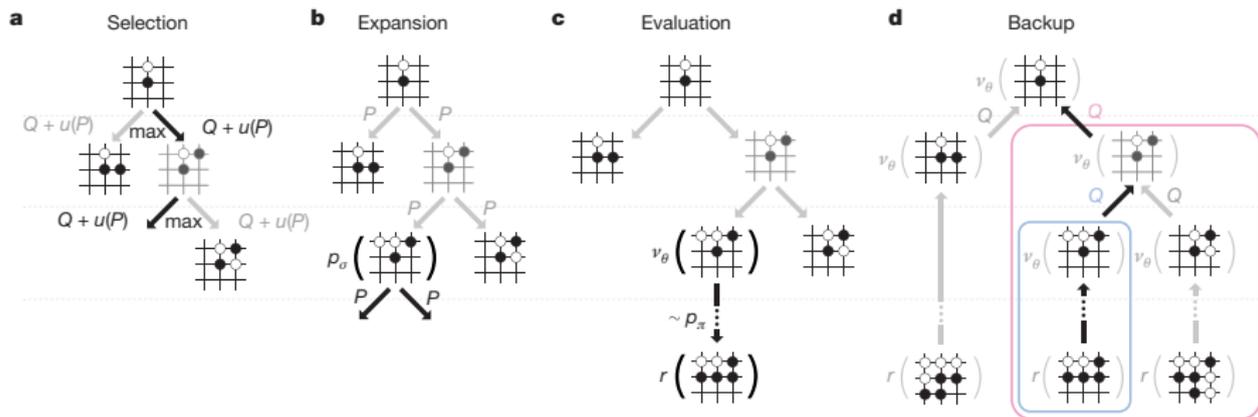


Third wave:
Contextual adaptation

(A DARPA Perspective on Artificial Intelligence,
<https://www.youtube.com/watch?v=-001G3tSYpU>)

Combining paradigms: AlphaGo and AlphaZero

AlphaGo & AlphaZero are combining symbolic and subsymbolic AI: heuristic search, reinforcement learning and deep neural networks.



Beating Lee Sedol (world-class Go player) in 2016 required approximately 2000 CPU cores and 300 GPUs, probably around 600kW. AlphaZero only 1-2 kW, but still specialised and more than most have available.

Silver et al.: Mastering the game of Go with deep neural networks and tree search. Nature, January 2016.

The computer that mastered Go

http://www2.compute.dtu.dk/~tobo/alphago_trimmed.mp4

(Nature Video: The Computer that Mastered Go (trimmed))

More Google DeepMind: Learning to play Atari

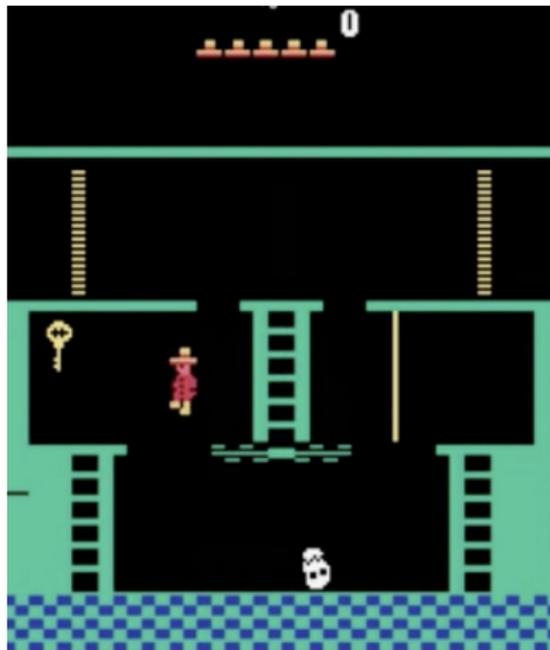
Combining **reinforcement learning** (Q-learning) and **deep neural networks**. Reported in Nature vol 518, 26 February 2015.

<http://www2.compute.dtu.dk/~tobo/DeepMind.MP4>

Easy or hard? Well-defined or not?



Google DeepMind 2017



Google DeepMind 2015

DeepMind playing Montezuma's Revenge

<http://www2.compute.dtu.dk/~tobo/MontezumasRevenge.mov>

The technology of Google's driverless cars

Symbolic and subsymbolic AI comes together.

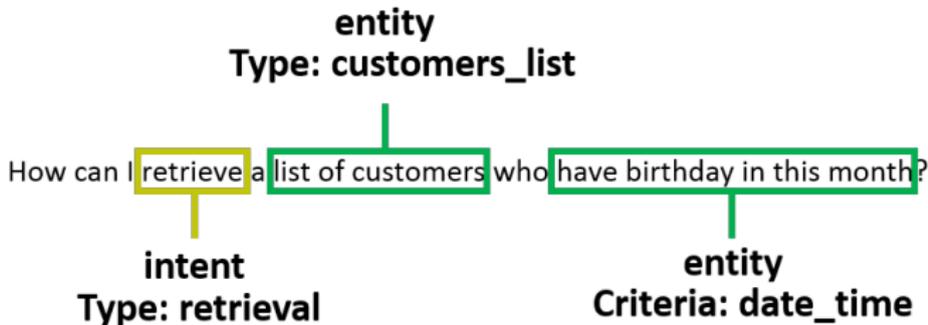
http://www2.compute.dtu.dk/~tobo/google_car_nosound.mp4

Explainable AI

- Cases with crashing driverless cars and racial AI have called for more **transparent** and **explainable** AI.
- Deep neural networks and genetic algorithms (subsymbolic) are naturally opaque.
- Manually hand-crafted rule-based systems (symbolic) are naturally transparent, but are normally unrealistic to craft.
- Best current bet is to **combine**: The output of learning is rules and explicit models that can be inspected, understood and modified by humans.

1st wave AI for natural language: chatbots from IBM Watson, Amazon Alexa, Wit.ai, etc.

Recognition of **intent** and **entities**.



“ book me a table for 5 at the Skates of Bay in Berkeley tonight ⚙

Intent	travel	⌵
wit/datetime	from 3/13/2015, 6:00:00 PM to 3/14/2015, 12:00:00 AM	×
wit/number:attendees	5	×
restaurant_name	Skates of Bay	⌵
wit/location:from	Berkeley	×
travel_type	restaurant	⌵

+ Add Entity Quick Add

✓ Validate ✕ Archive

Amelia for SEB bank customers (since Jan 2017)

Dialogue between me and Amelia (Aida) at the home page of SEB, 29 May 2017:

Thomas: Hur öppnar jeg ett konto?

Aida (Amelia): Är du kund hos oss sedan tidigare?

Thomas: Nej.

Aida: För at skaffa Privatkonto behöver du vara kund i SEB. Vill du at jeg guidar dig?

Thomas: Ja tak.

Aida: För att kunne hitta den enklaste lösningen för dig behöver jeg ställa några frågor. Är du över 18 år?

Thomas: Ja.

Aida: Kan du legitimera dig med mobilt BankID?

Thomas: Jag är från Danmark.

Aida: Jag förstår inte riktigt vad du vill ha hjälp med, men formulera dig gärna på ett annat sätt. Vilket är ditt svar?

SEB Aida

<http://www2.compute.dtu.dk/~tobo/SIRIkomSEB.mp4>

2nd wave AI for natural language: OpenAI's language model GPT-2, February 2019

A **language model** has been trained on massive amounts of text in order to be able to statistically predict the next word in a sentence. Using recurrent neural networks (subsymbolic).

Sentence:

I'd like a nice hot cup of

Options:

45.5% **coffee**
32.6% **tea**
2.2% **hot**
1.4% **jo**
1.0% **espresso**
0.8% **cocoa**
0.6% **water**
0.6% **chocolate**
0.5% **milk**
0.3% **ice**

Sentence:

I'd like a nice hot cup of skimmed

Options:

82.8% **milk**
2.7% **coffee**
1.5% **egg**
1.3% **water**
1.3% **tea**
0.7% **butter**
0.7% **-**
0.7% **lemon**
0.4% **sugar**
0.3% **cocoa**

(Screenshots from the public (less precise) model, <https://gpt2.apps.openai.org>)

GPT-2 demo: Input taken from BBC article, 28 November 2019

US President Donald Trump has signed into law a bill that supports pro-democracy protesters in Hong Kong. The Human Rights and Democracy Act mandates an annual review, to check if Hong Kong has enough autonomy to justify its special status with the US.

Mr Trump said he signed the law "out of respect for President Xi [Jinping], China, and the people of Hong Kong". China's foreign ministry threatened "counter measures" if the US continued "going down the wrong path". Mr Trump is currently seeking a deal with China, in order to end a trade war between the two countries. "The US has been disregarding facts and distorting truth," a Chinese foreign ministry statement said. "It openly backed violent criminals who rampantly smashed facilities, set fire, assaulted innocent civilians, trampled on the rule of law, and jeopardized social order."

Hong Kong's government also reacted, saying the American bill would send the wrong signal and would not help to ease the situation. But a key activist in the Hong Kong protest movement, Joshua Wong, said the US law was a "remarkable achievement" for "all Hongkongers".

Mr Trump had previously been non-committal about whether he would sign the bill, saying he was "with" Hong Kong but also that Mr Xi was "an incredible guy". However, the bill had widespread congressional support, which meant that even if he vetoed it, lawmakers could potentially have voted to overturn his decision. The president also signed a second bill, which bans the export of crowd-control munitions to the police in Hong Kong - including tear gas, rubber bullets and stun guns. "[The bills] are being enacted in the hope that leaders and representatives of China and Hong Kong will be able to amicably settle their differences, leading to long-term peace and prosperity for all," Mr Trump said.

GPT-2 demo: Output produced by model

China "does not wish to see instability or war in any region of the world," said the White House. Mr Trump's visit to China comes at the end of a month that has seen a wave of protests across the South China Sea. He will be meeting Mr Xi at the Great Hall of the People on Sunday - the site of the last time the leaders met. The two presidents will also discuss trade and North Korea.

President Donald Trump speaks on the steps of the White House during the weekly policy luncheon in Washington, DC

The US President will be meeting his Chinese counterpart Xi Jinping at the Great Hall of the People in Beijing.

Mr Trump's visit comes at the end of a month that has seen a wave of protests across the South China Sea.

The White House also issued a statement saying that Mr Trump would urge the "two nations to advance a new model of major-power relations" and work together to advance security, prosperity, and peace in the region. The statement added that Mr Xi had expressed his country's willingness to work with the US on regional issues, including North Korea. The US President is on a 10-day trip that will see him make stops in Saudi Arabia, Israel and the Vatican. He will then go to the Middle East before returning to the United States on Friday.

The White House statement said that Mr Trump would urge the "two nations to advance a new model of major-power relations" and work together to advance security, prosperity, and peace in the region

Mr Trump will also visit the Vatican and attend the traditional Nato summit in Brussels.

The Chinese capital was last week hit by massive protests after Mr Trump criticised the Beijing government on Twitter. As many as 90,000 people took to the streets of Hong Kong to demand more freedom and democracy.