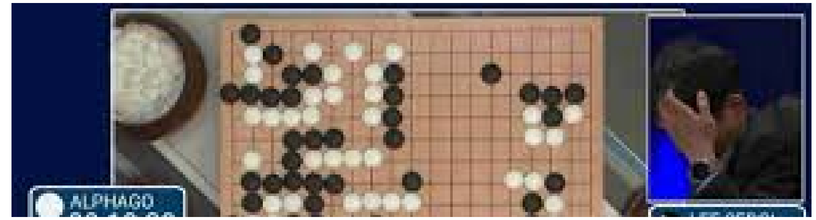

Microsoft creates AI that can read a document and answer questions about it as well as a person

January 15, 2018 | Allison Linn

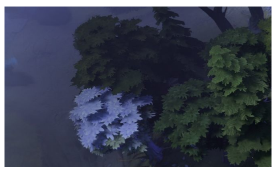

Microsoft researchers achieve new conversational speech recognition milestone

August 20, 2017 | By Xuedong

June 24, 2014

DeepFace: Closing Performance in Fa

Conference on Computer Vision and Pattern Recognition (CVPR)

By: **Yaniv Taigman**, Ming Yang, **Marc'Aurelio Ranzato**, **Lior Wolf**

Abstract

In modern face recognition, the conventional pipeline co
classify. We revisit both the alignment step and the repre
modeling in order to apply a piecewise affine transforma
layer deep neural network. This deep network involves m
locally connected layers without weight sharing, rather t
trained it on the largest facial dataset to-date, an identity
belonging to more than 4,000 identities.

If you think AI will never replace radiologists—you may want to think again

May 14, 2018 | Michael Walter | Artificial Intelligence

It's one of the most frequently discussed questions in radiology today: What kind of long-term impact will artificial intelligence (AI) have on radiologists?

Robert Schier, MD, a radiologist for RadNet, shared his own thoughts on the topic in a new commentary published by the *Journal of the American College of Radiology*—and he's not quite as optimistic as some of his colleagues throughout the industry.

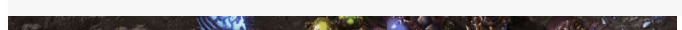

- It is generally not hard to motivate AI these days. There have been some substantial success stories. A lot of the triumphs have been in **games**, such as Jeopardy! (IBM Watson, 2011), Go (DeepMind's AlphaGo, 2016), Dota 2 (OpenAI, 2019), Poker (CMU and Facebook, 2019).
- On non-game tasks, we also have systems that achieve strong performance on reading comprehension, speech recognition, face recognition, and medical imaging **benchmarks**.
- Unlike games, however, where the game is the full problem, good performance on a benchmark does not necessarily translate to good performance on the actual task in the wild. Just because you ace an exam doesn't necessarily mean you have perfect understanding or know how to apply that knowledge to real problems.
- So, while promising, not all of these results translate to real-world applications

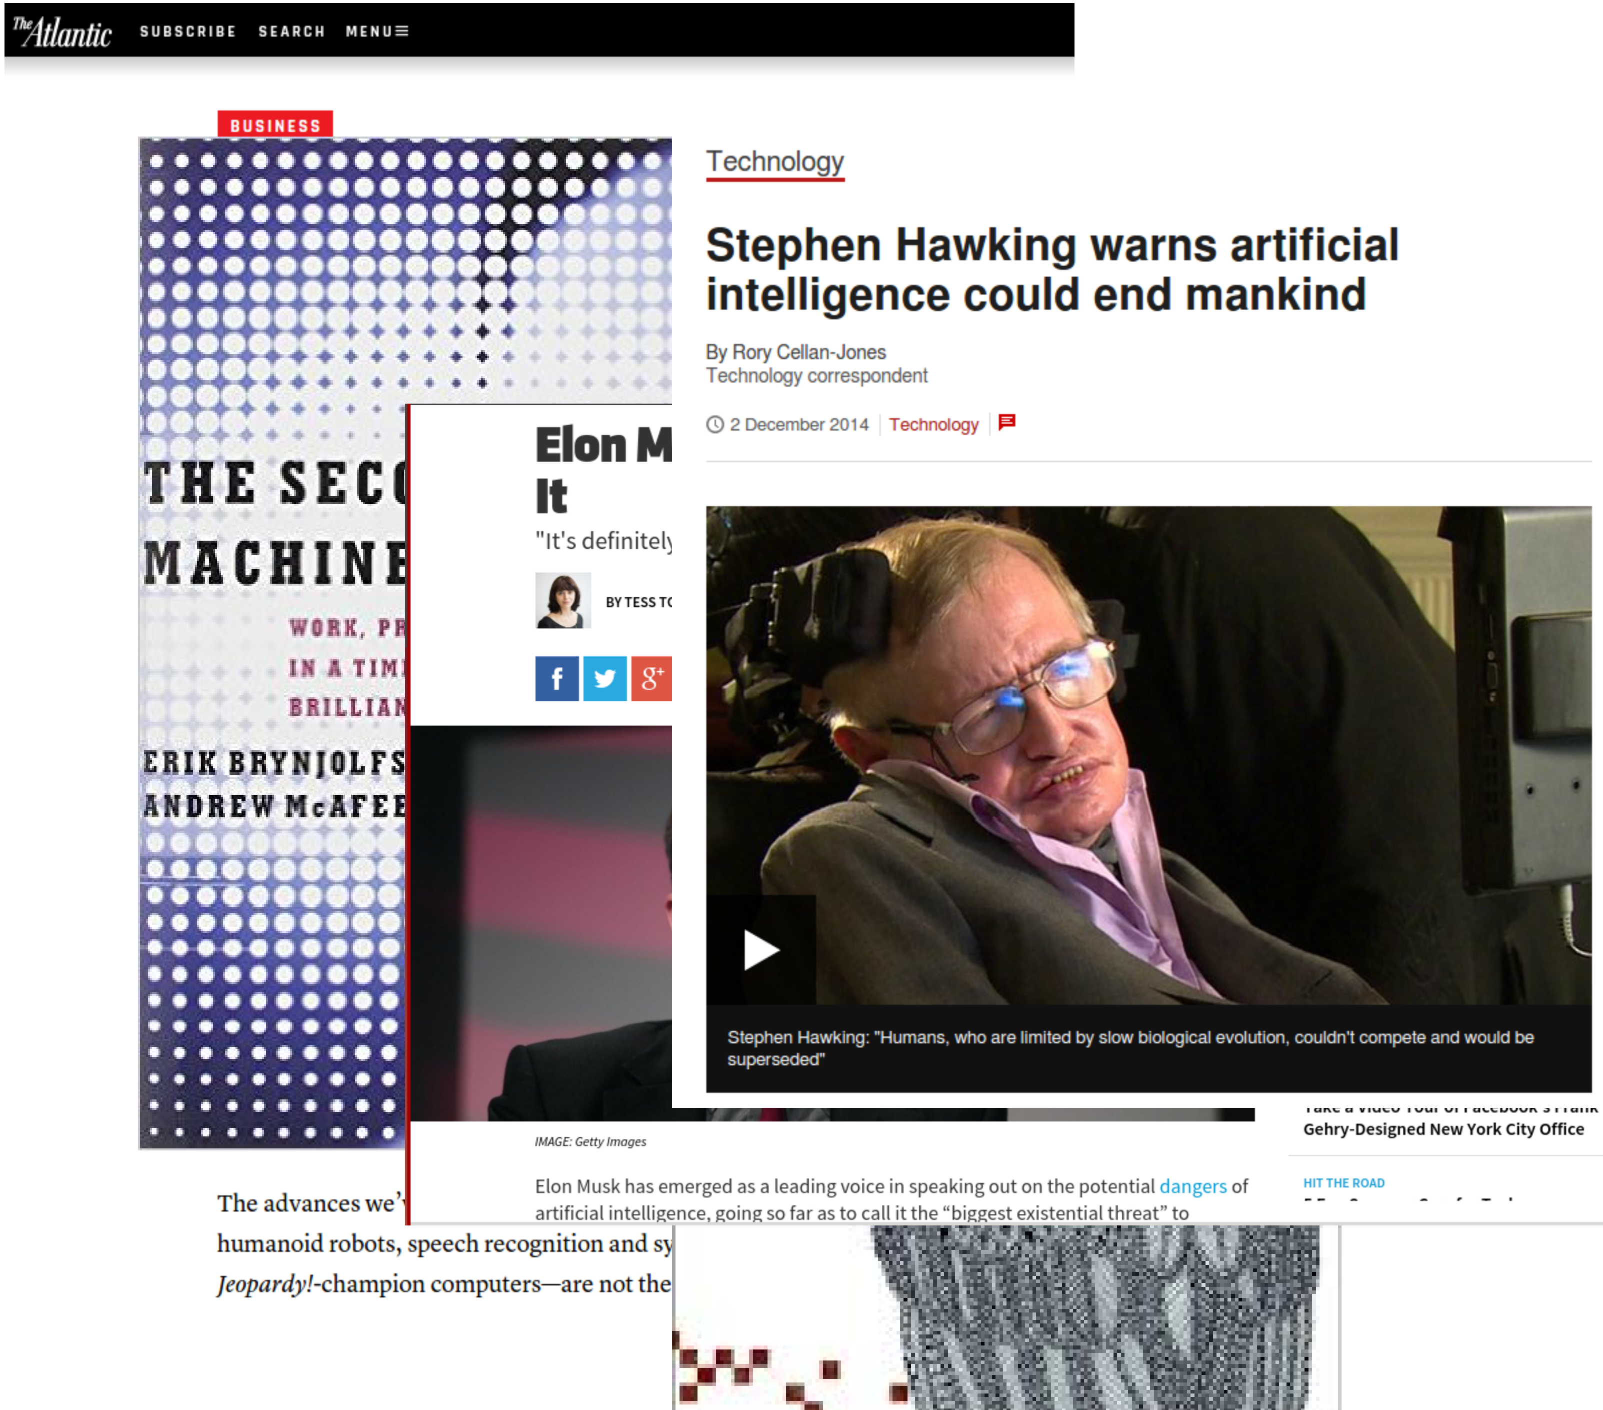
The Atlantic
SUBSCRIBE
SEARCH
MENU
BUSINESS
THE SECO
MACHINE
WORK, PR
IN A TIM
BRILLIAN
ERIK BRYNJOLFS
ANDREW McAFEE
The advances we'
humanoid robots, speech recognition and sy
Jeopardy!-champion computers—are not the
Elon M
It
"It's definitely
BY TESS TC
IMAGE: Getty Images
Elon Musk has emerged as a leading voice in speaking out on the potential dangers of
artificial intelligence, going so far as to call it the "biggest existential threat" to
Gehry-Designed New York City Office
HIT THE ROAD
Technology
Stephen Hawking warns artificial intelligence could end mankind
By Rory Cellan-Jones
Technology correspondent
2 December 2014
Technology
Stephen Hawking: "Humans, who are limited by slow biological evolution, couldn't compete and would be superseded"

- From the non-scientific community, we also see speculation about the future: that it will bring about sweeping societal change due to automation, resulting in massive job loss, not unlike the industrial revolution, or that AI could even surpass human-level intelligence and seek to take control.
- While these are extreme views, there is no doubt that AI is and will continue to be transformational. We still don't know exactly what that transformation will look like.

1956

Birth of AI

1956: Workshop at Dartmouth College; attendees: John McCarthy, Marvin Minsky, Claude Shannon, etc.

Aim for **general principles**:

Every aspect of learning or any other feature of intelligence can be so precisely described that a machine can be made to simulate it.

- How did we get here? The name **artifical intelligence** goes back to a summer in 1956. John McCarthy, who was then at MIT but later founded the Stanford AI lab, organized a workshop at Dartmouth College with the leading thinkers of the time, and set out a very bold proposal...to build a system that could do it **all**.

Birth of AI, early successes

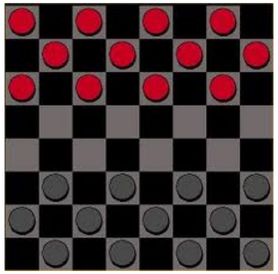

Checkers (1952): Samuel's program learned weights and played at strong amateur level

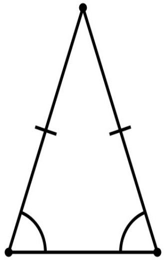

Problem solving (1955): Newell & Simon's Logic Theorist: prove theorems in Principia Mathematica using search + heuristics; later, General Problem Solver (GPS)

- While they did not solve it all, there were a lot of **interesting programs** that were created: programs that could play checkers at a strong amateur level, programs that could prove theorems.
- For one theorem Newell and Simon's Logical Theorist actually found a proof that was more elegant than what a human came up with. They actually tried to publish a paper on it but it got rejected because it was not a new theorem; perhaps they failed to realize that the third author was a computer program.
- From the beginning, people like John McCarthy sought **generality**, thinking of how commonsense reasoning could be encoded in logic. Newell and Simon's General Problem Solver promised to solve any problem (which could be suitably encoded in logic).

Overwhelming optimism...

Machines will be capable, within twenty years, of doing any work a man can do. —Herbert Simon

Within 10 years the problems of artificial intelligence will be substantially solved. —Marvin Minsky

I visualize a time when we will be to robots what dogs are to humans, and I'm rooting for the machines. —Claude Shannon

- It was a time of high optimism, with all the leaders of the field, all impressive thinkers, predicting that AI would be "solved" in a matter of years.

...underwhelming results

Example: machine translation

The spirit is willing but the flesh is weak.

↓

(Russian)

↓

The vodka is good but the meat is rotten.

1966: ALPAC report cut off government funding for MT, first AI winter

- Despite some successes, certain tasks such as machine translation were complete failures, which lead to the cutting of funding and the first AI winter.

Implications of early era

Problems:

- **Limited computation**: search space grew exponentially, outpacing hardware ($100! \approx 10^{157} > 10^{80}$)
- **Limited information**: complexity of AI problems (number of words, objects, concepts in the world)

Contributions:

- Lisp, garbage collection, time-sharing (John McCarthy)
- Key paradigm: separate **modeling** and **inference**

- What went wrong? It turns out that the real world is very complex and most AI problems require a lot of **compute** and **data**.
- The hardware at the time was simply too limited both compared to the human brain and computers available now. Also, casting problems as general logical reasoning meant that the approaches fell prey to the exponential search space, which no possible amount of compute could really fix.
- Even if you had infinite compute, AI would not be solved. There are simply too many words, objects, and concepts in the world, and this information has to be somehow encoded in the AI system.
- Though AI was not solved, a few generally useful technologies came out of the effort, such as Lisp (~~still the world's most advanced programming language in a sense~~).
- One particularly powerful paradigm is the separation between what you want to compute (modeling) and how to compute it (inference).

Knowledge-based systems (70-80s)

Expert systems: elicit specific domain knowledge from experts in form of rules:

```
if [premises] then [conclusion]
```

- In the seventies and eighties, AI researchers looked to knowledge as a way to combat both the limited computation and information problems. If we could only figure out a way to encode prior knowledge in these systems, then they would have the necessary information and also have to do less compute.

Knowledge-based systems (70-80s)

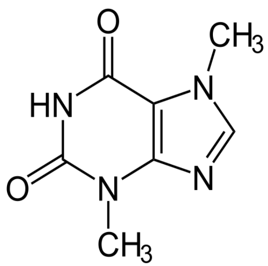

DENDRAL: infer molecular structure from mass spectrometry

MYCIN: diagnose blood infections, recommend antibiotics

XCON: convert customer orders into parts specification; save DEC $\$40$ million a year by 1986

- Instead of the solve-it-all optimism from the 1950s, researchers focused on building narrow practical systems in targeted domains. These became known as **expert systems**.

Knowledge-based systems

Contributions:

- First **real application** that impacted industry
- Knowledge helped curb the exponential growth

Problems:

- Knowledge is not deterministic rules, need to model **uncertainty**
- Requires considerable **manual effort** to create rules, hard to maintain

1987: Collapse of Lisp machines and second AI winter

- This was the first time AI had a measurable impact on industry. However, the technology ran into limitations and failed to scale up to more complex problems. Due to plenty of overpromising and underdelivering, the field collapsed again.
- We know that this is not the end of the AI story, but actually it is not the beginning. There is another thread for which we need to go back to 1943.

1943

Artificial neural networks

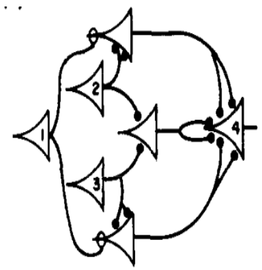

1943: introduced artificial neural networks, connect neural circuitry and logic (McCulloch/Pitts)

1969: Perceptrons book showed that linear models could not solve XOR, killed neural nets research (Minsky/Papert)

- Much of AI's history was dominated by the logical tradition, but there was another smaller camp, grounded in neural networks inspired by the brain.
- (Artificial) neural networks were introduced by a famous paper by McCulloch and Pitts, who devised a simple mathematical model and showed how it could be be used to compute arbitrary logical functions.
- Much of the early work was on understanding the mathematical properties of these networks, since computers were too weak to do anything interesting.
- In 1969, a book was published that explored many mathematical properties of Perceptrons (linear models) and showed that they could not solve some simple problems such as XOR. Even though this result says nothing about the capabilities of deeper networks, the book is largely credited with the demise of neural networks research, and the continued rise of logical AI.

Training networks

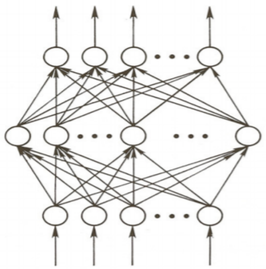

1986: popularization of backpropagation for training multi-layer networks (Rumelhardt, Hinton, Williams)

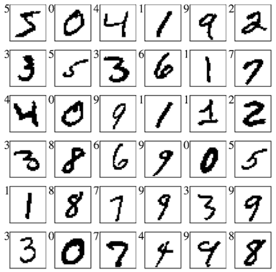

1989: applied convolutional neural networks to recognizing handwritten digits for USPS (LeCun)

- In the 1980s, there was a renewed interest in neural networks. Backpropagation was rediscovered and popularized as a way to actually train deep neural networks, and Yann LeCun built a system based on convolutional neural networks to recognize handwritten digits. This was one of the first successful uses of neural networks, which was then deployed by the USPS to recognize zip codes.

Deep learning

AlexNet (2012): huge gains in object recognition; transformed computer vision community overnight

AlphaGo (2016): deep reinforcement learning, defeat world champion Lee Sedol

- The real break for neural networks came in the 2010s. With the rise of compute (notably GPUs) and large datasets such as ImageNet (2009), the time was ripe for the world to take note of neural networks.
- AlexNet was a pivotal system that showed the promise of deep convolutional networks on ImageNet, the benchmark created by the computer vision community who was at the time still skeptical of deep learning. Many other success stories in speech recognition and machine translation followed.

Two intellectual traditions

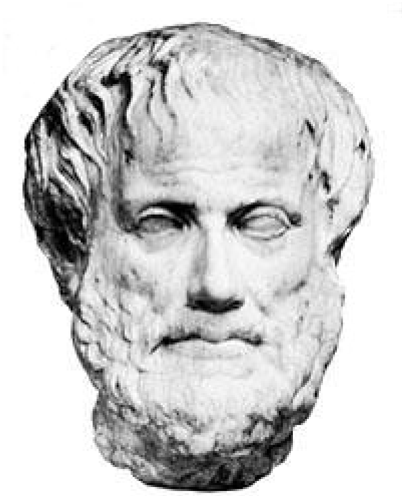

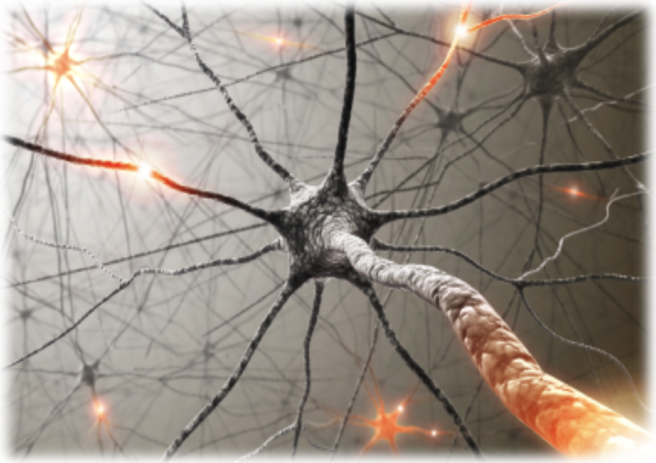

- AI has always swung back and forth between the two
- Deep philosphical differences, but deeper connections (McCulloch/Pitts, AlphaGo)?

- Reflecting back on the past of AI, there have been two intellectual traditions that have dominated the scene: one rooted in logic and one rooted in neuroscience (at least initially). This debate is paralleled in cognitive science with connectionism and computationalism.
- While there are deep philosophical differences, perhaps there are deeper connections.
- For example, McCulloch and Pitts' work from 1943 can be viewed as the root of deep learning, but that paper is mostly about how to implement logical operations.
- The game of Go (and indeed, many games) can be perfectly characterized by a set of simple logic rules. At the same time, the most successful systems (AlphaGo) do not tackle the problem directly using logic, but appeal to the fuzzier world of artificial neural networks.

A melting pot

- Bayes rule (Bayes, 1763) from **probability**
- Least squares regression (Gauss, 1795) from **astronomy**
- First-order logic (Frege, 1893) from **logic**
- Maximum likelihood (Fisher, 1922) from **statistics**
- Artificial neural networks (McCulloch/Pitts, 1943) from **neuroscience**
- Minimax games (von Neumann, 1944) from **economics**
- Stochastic gradient descent (Robbins/Monro, 1951) from **optimization**
- Uniform cost search (Dijkstra, 1956) from **algorithms**
- Value iteration (Bellman, 1957) from **control theory**

- Of course, any story is incomplete.
- In fact, for much of the 1990s and 2000s, neural networks were not popular in the machine learning community, and the field was dominated more by techniques such as Support Vector Machines (SVMs) inspired by statistical theory.
- The fuller picture is that the modern world of AI is more like New York City—it is a melting pot that has drawn from many different fields ranging from statistics, algorithms, economics, etc.
- And often it is the new connections between these fields that are made and their application to important real-world problems that makes working on AI so rewarding.

Two views of AI

AI agents: how can we create intelligence?

AI tools: how can we benefit society?

- There are two ways to look at AI philosophically.
- The first is the science and engineering of building "intelligent" agents. The inspiration of what constitutes intelligence comes from the types of capabilities that humans possess: the ability to perceive a very complex world and make enough sense of it to be able to manipulate it.
- The second views AI as a set of tools. We are simply trying to solve problems in the world, and techniques developed by the AI community happen to be useful for that, but these problems are not ones that humans necessarily do well on natively.
- However, both views boil down to many of the same day-to-day activities (e.g., collecting data and optimizing a training objective), the philosophical differences do change the way AI researchers approach and talk about their work. Moreover, the conflation of these two views can generate a lot of confusion.

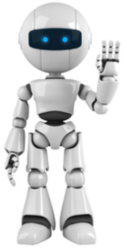

AI agents...

An intelligent agent

Perception | Robotics | Language

Knowledge | Reasoning | Learning

- The starting point for the agent-based view is ourselves.
- As humans, we have to be able to perceive the world (computer vision), perform actions in it (robotics), and communicate with other agents (language).
- We also have knowledge about the world (from procedural knowledge like how to ride a bike, to declarative knowledge like remembering the capital of France), and using this knowledge we can draw inferences and make decisions (reasoning).
- Finally, we learn and adapt over time. We are born with none of the skills that we possess as adults, but rather the capacity to acquire them. Indeed machine learning has become the primary driver of many of the AI applications we see today.

Are we there yet?

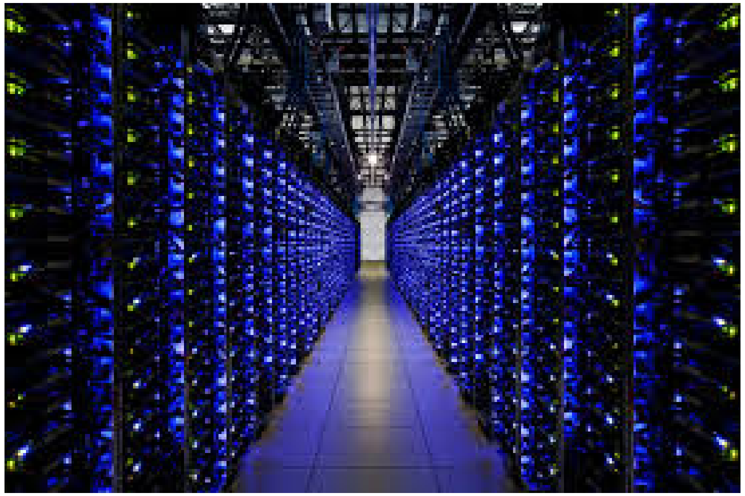

Machines: narrow tasks, millions of examples

Humans: diverse tasks, very few examples

- The AI agents view is an inspiring quest to undercover the mysteries of intelligence and tackle the tasks that humans are good at. While there has been a lot of progress, we still have a long way to go along some dimensions: for example, the ability to learn quickly from few examples or the ability to perform commonsense reasoning.
- There is still a huge gap between the regimes that humans and machines operate in. For example, AlphaGo learned from 19.6 million games, but can only do one thing: play Go. Humans on the other hand learn from a much wider set of experiences, and can do many things.

 *AI tools...*

- The other view of AI is less about re-creating the capabilities that humans have, and more about how to benefit humans.
- Even the current level of technology is already being deployed widely in practice, and many of these settings are often not particularly human-like (targeted advertising, news or product recommendation, web search, supply chain management, etc.)

Predicting poverty

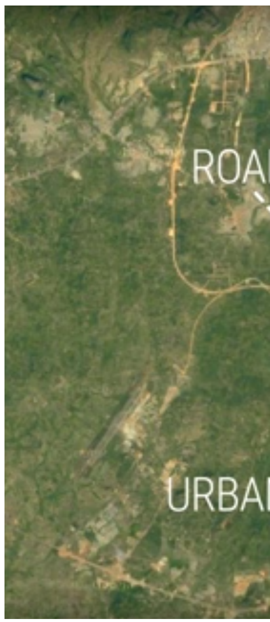

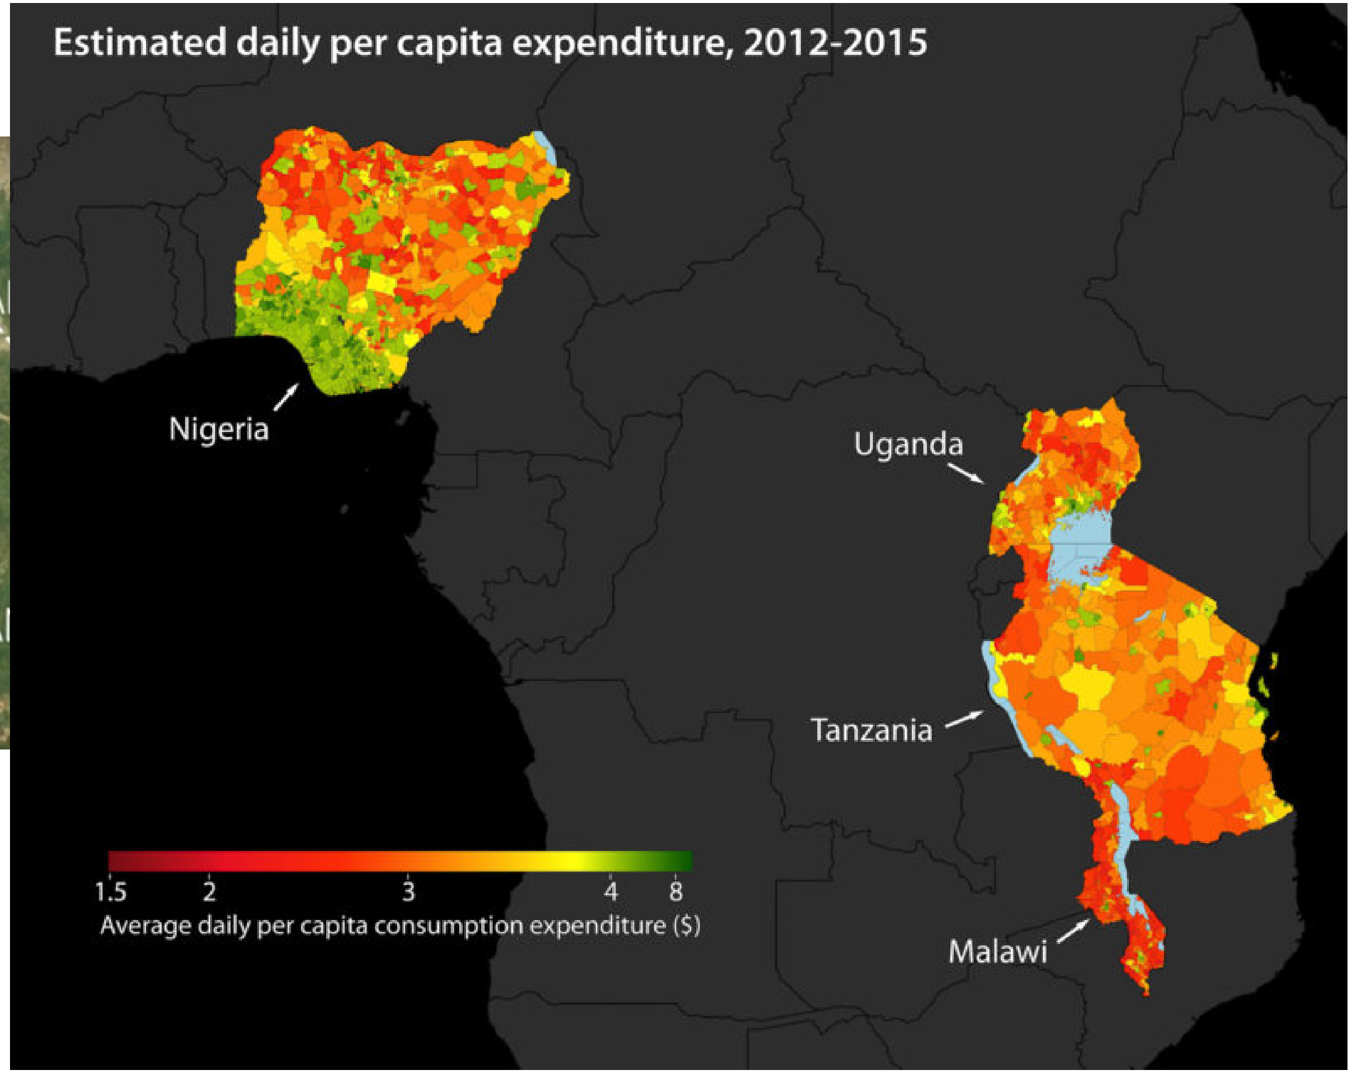

- Computer vision techniques, used to recognize objects, can also be used to tackle social problems. Poverty is a huge problem, and even identifying the areas of need is difficult due to the difficulty in getting reliable survey data. Recent work has shown that one can take satellite images (which are readily available) and predict various poverty indicators.

Saving energy by cooling datacenters

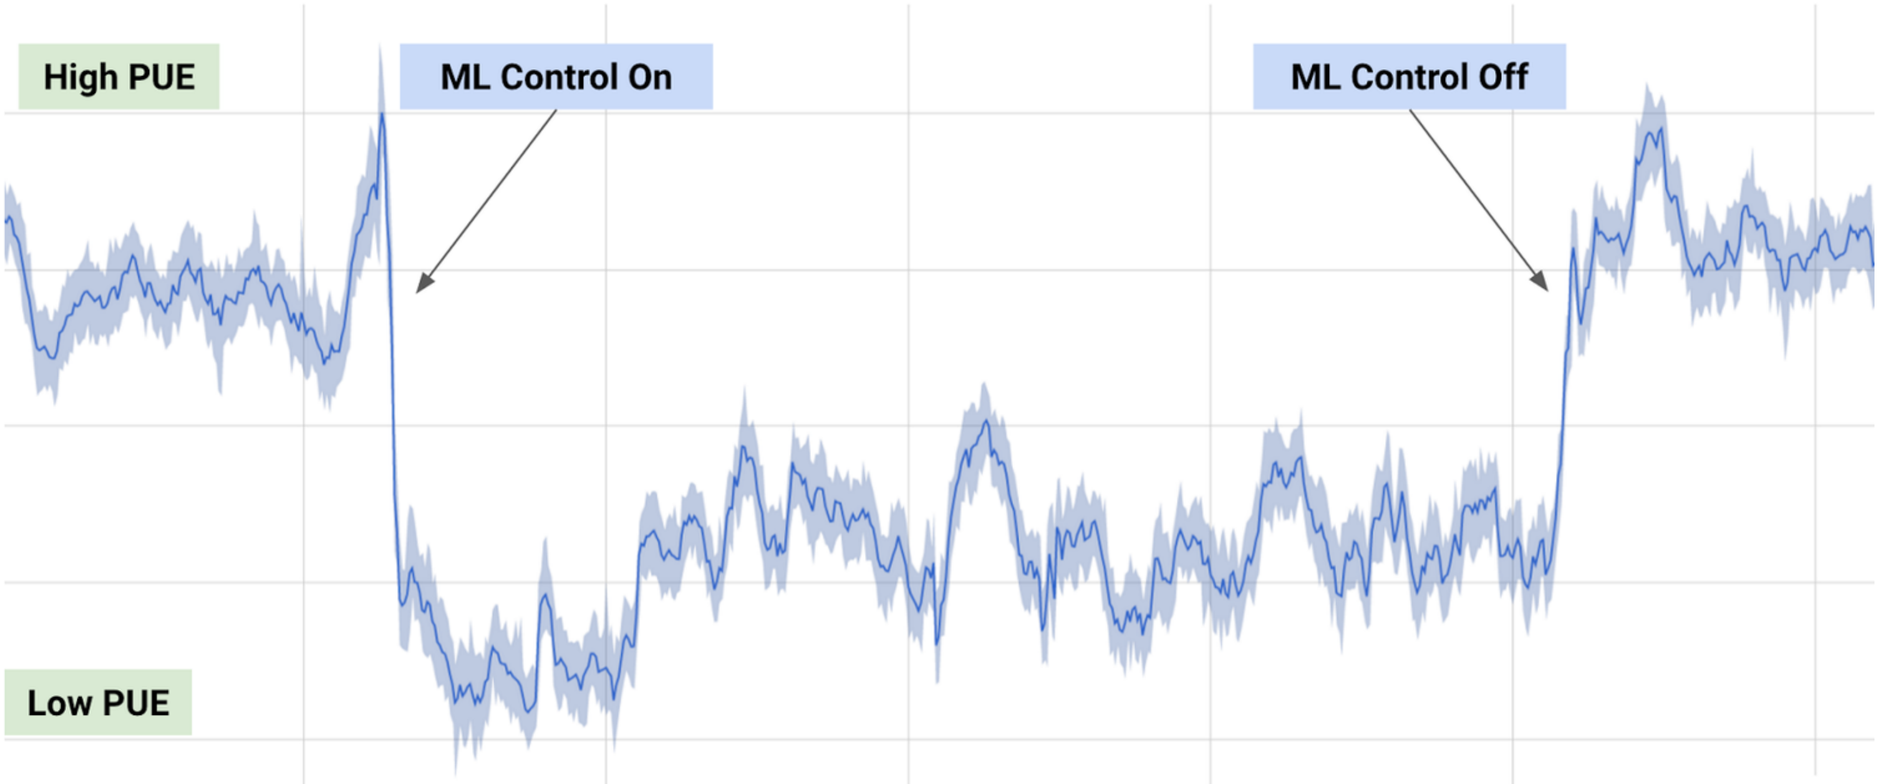

- Machine learning can also be used to optimize the energy efficiency of datacenters which, given the hunger for compute these days, makes a big difference. Some recent work from DeepMind shows how to significantly reduce Google's energy footprint by using machine learning to predict the power usage effectiveness from sensor measurements such as pump speeds, and use that to drive recommendations.

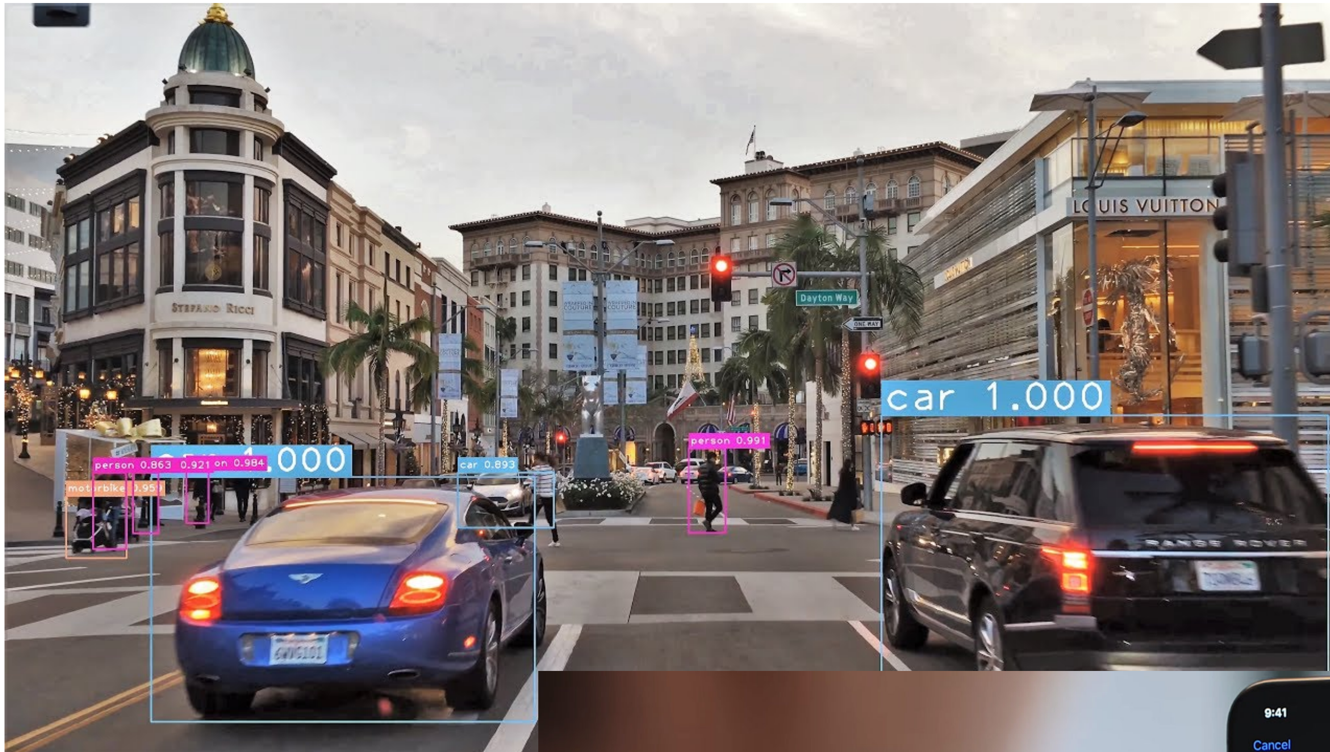
car 1.000
person 0.991
car 0.893
LOUIS VUITTON
Dayton Way

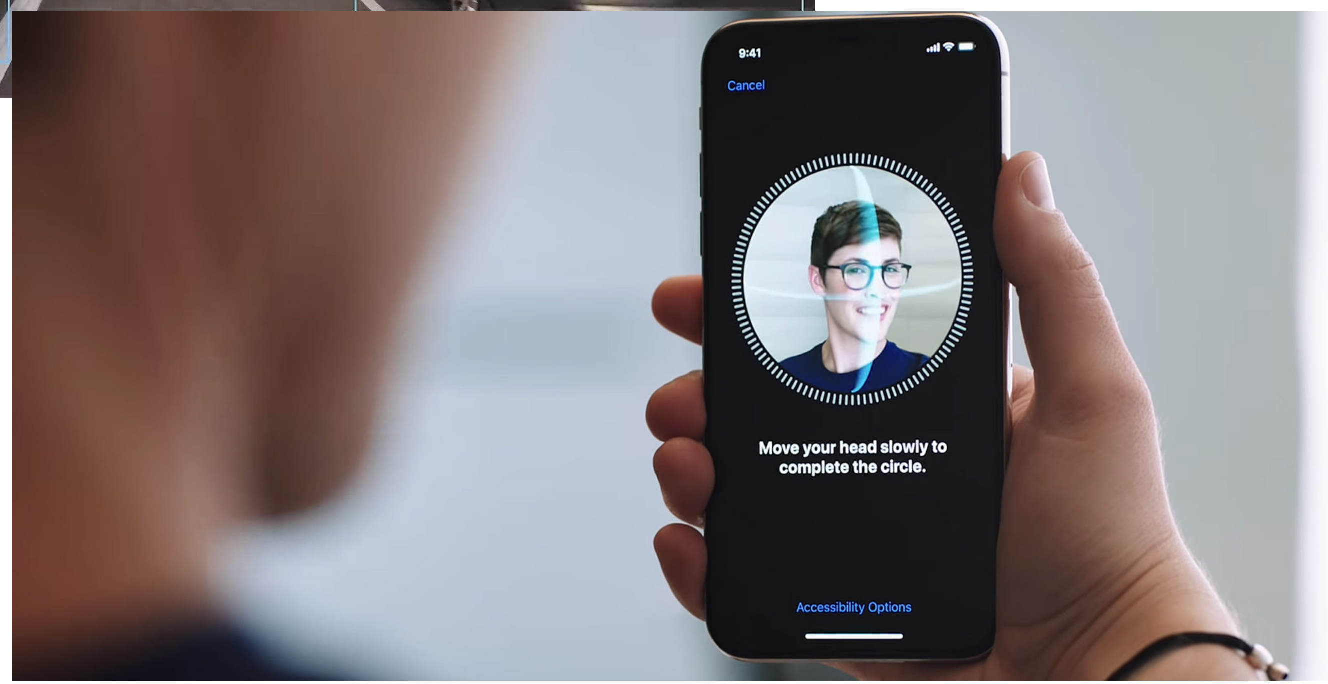
9:41
Cancel
Move your head slowly to complete the circle.
Accessibility Options

Security

[Evtimov+ 2017]

[Sharif+ 2016]

- Other applications such as self-driving cars and authentication have high-stakes, where errors could be much more damaging than getting the wrong movie recommendation. These applications present a set of security concerns.
- One can generate so-called **adversarial examples**, where by putting stickers on a stop sign, one can trick a computer vision system into mis-classifying it as a speed limit sign. You can also purchase special glasses that fool a system into thinking that you're a celebrity.

Bias in machine translation

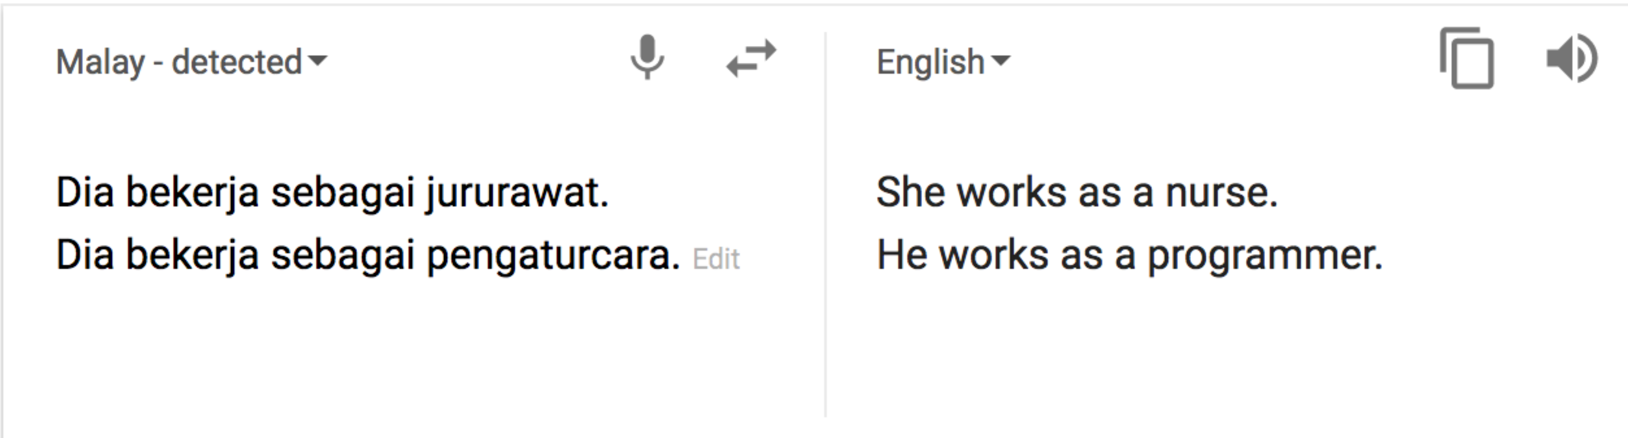

society $\Rightarrow$ data $\Rightarrow$ predictions

- A more subtle case is the issue of bias. One might naively think that since machine learning algorithms are based on mathematical principles, they are somehow objective. However, machine learning predictions come from the training data, and the training data comes from society, so any biases in society are reflected in the data and propagated to predictions. The issue of bias is a real concern when machine learning is used to decide whether an individual should receive a loan or get a job.
- Unfortunately, the problem of fairness and bias is as much of a philosophical one as it is a technical one. There is no obvious "right thing to do", and it has even been shown mathematically that it is impossible for a classifier to satisfy three reasonable fairness criteria (Kleinberg et al., 2016).

Summary so far

- AI agents: achieving human-level intelligence, still very far (e.g., generalize from few examples)

- AI tools: need to think carefully about real-world consequences (e.g., security, biases)