Why we will not bow to our

Computer Overlords
(just yet)

Grace Tang, Data Scientist
UBER

Many of you will recognise that title of this talk is inspired by the episode of Jeopardy where two of the most successful human Jeopardy contestants competed against IBM Watson (Ken Jennings and Brad Rutter, 2008).
Watson, backed by the power of the internet, obviously won, to which one of the humans reacted with the now famous phrase: "I for one, welcome our new computer overlords"
This is one of the more well-known events that perfectly captures our inherent distrust of AI, and the threat that its rapid development poses.

(gif source:
http://www.davidirons.com/thoughts/2016/2/4/the-greatest-ad-ever-then-now)

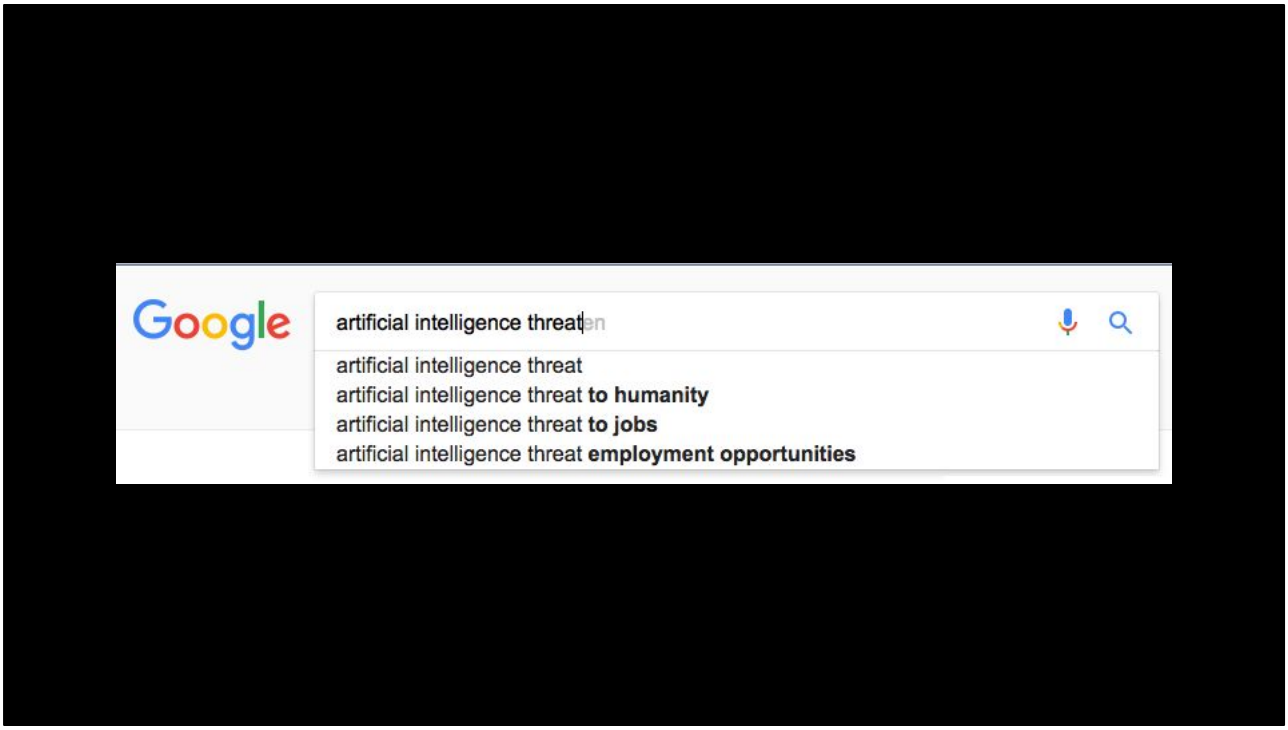

This fear is widespread. If you google the phrase “AI threatens [blank]”, we get the following autocompleted phrases:
AI threatens humanity, and AI threatens jobs.

AI = EVIL ??

Will AI make humans obsolete?

So let’s talk about those things today.

- Is AI actually evil?
- And willl AI replace our jobs and make us obsolete?

AI in Pop Culture

To start, let's first take a look at how AI is represented in popular culture.
What are some examples of fictional AI that come to mind?

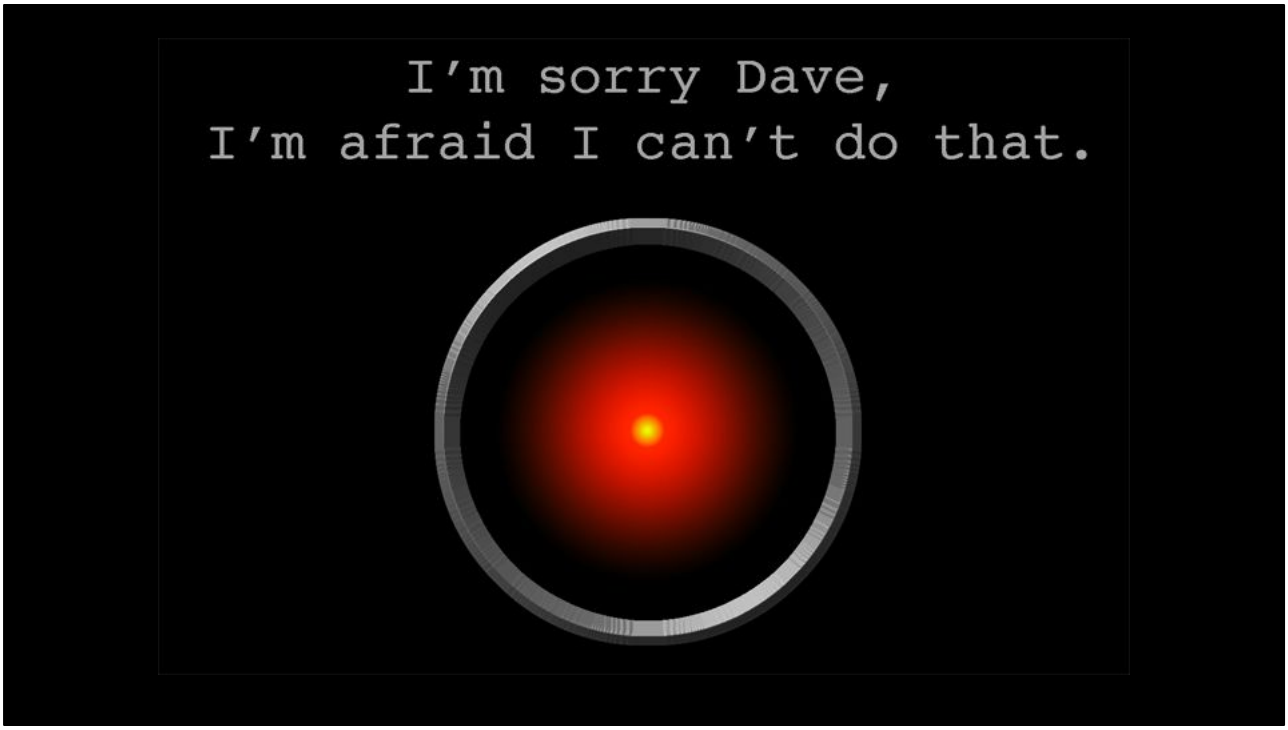

One classic example is HAL 9000, from 2001: A space odyssey (1968)… HAL kills its crew after they plan to disconnect it

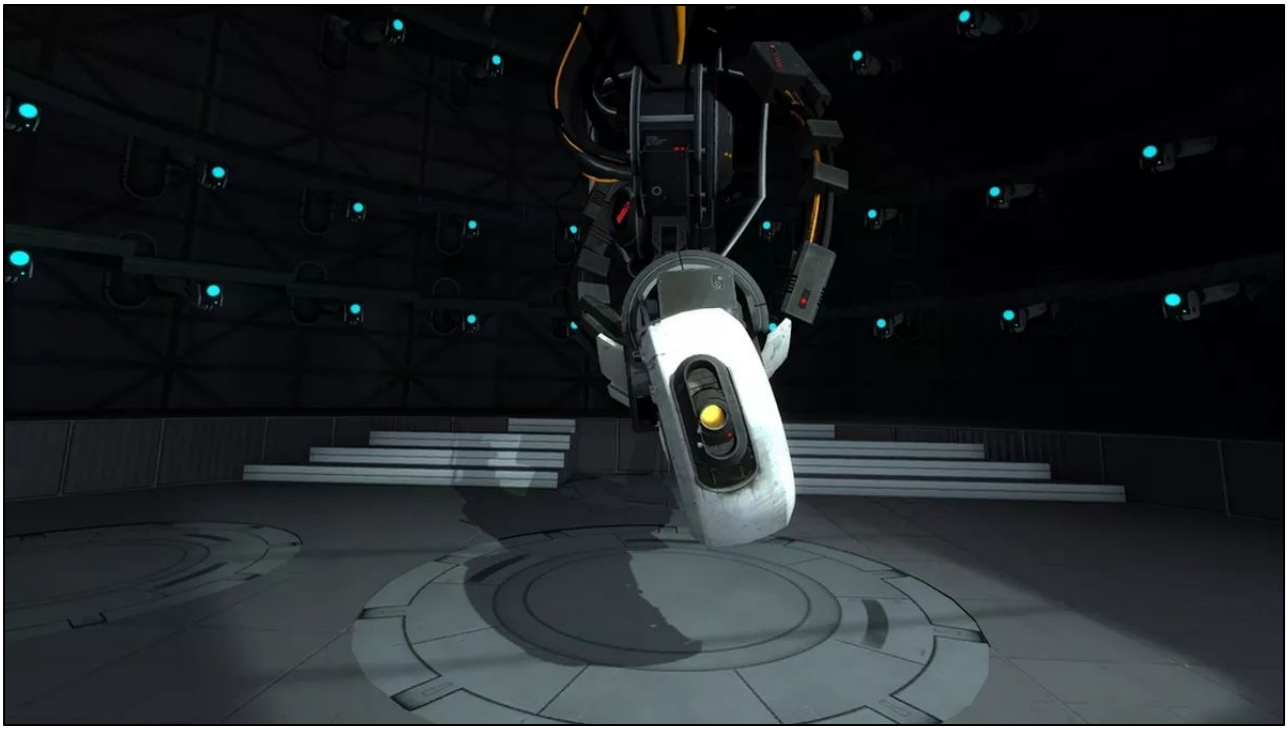

Then there’s Glados from Portal, A lab assistant robot who kills the human scientists who mistreat “her”, then runs experiments on innocent humans

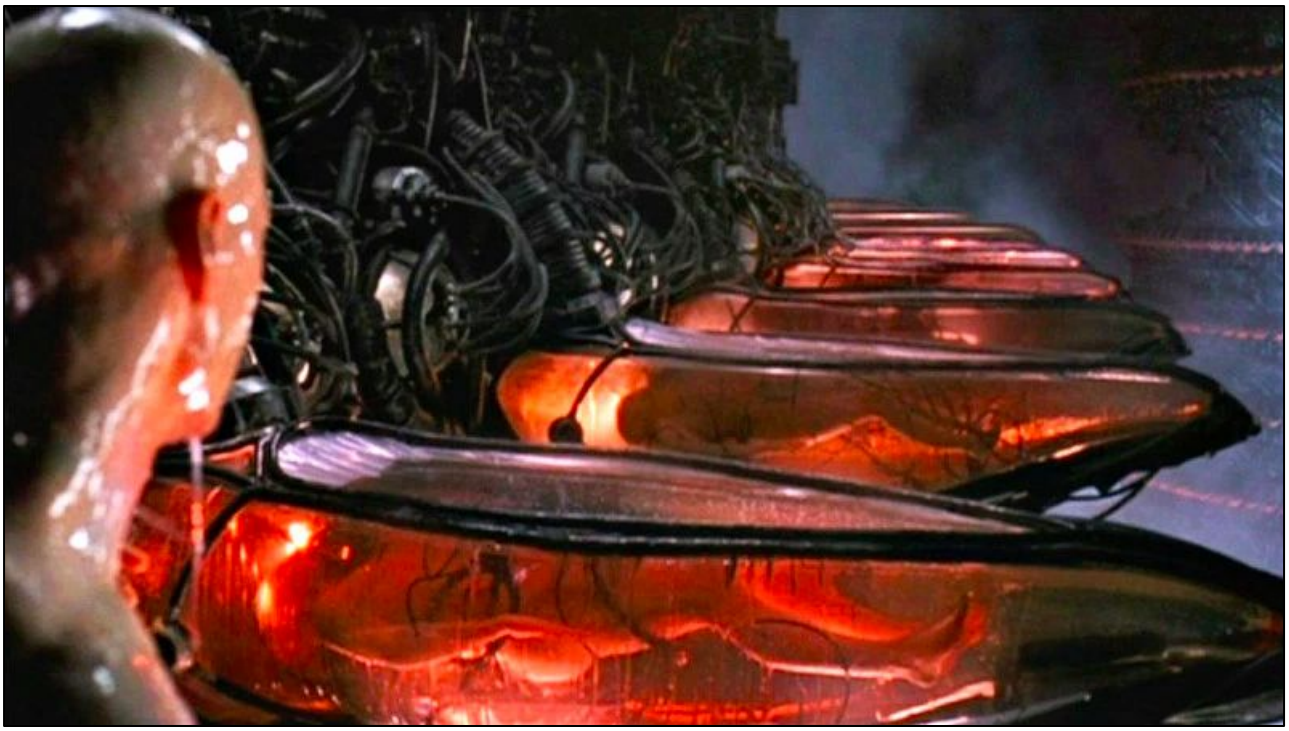

If we’re not lab rats, we’re batteries. In the matrix, humans are used by intelligent robots as a power source*

AI is depicted as being very sinister!

---

*originally, the plot involved robots using human minds as a literal neural network

AI in the Media

What about the media?

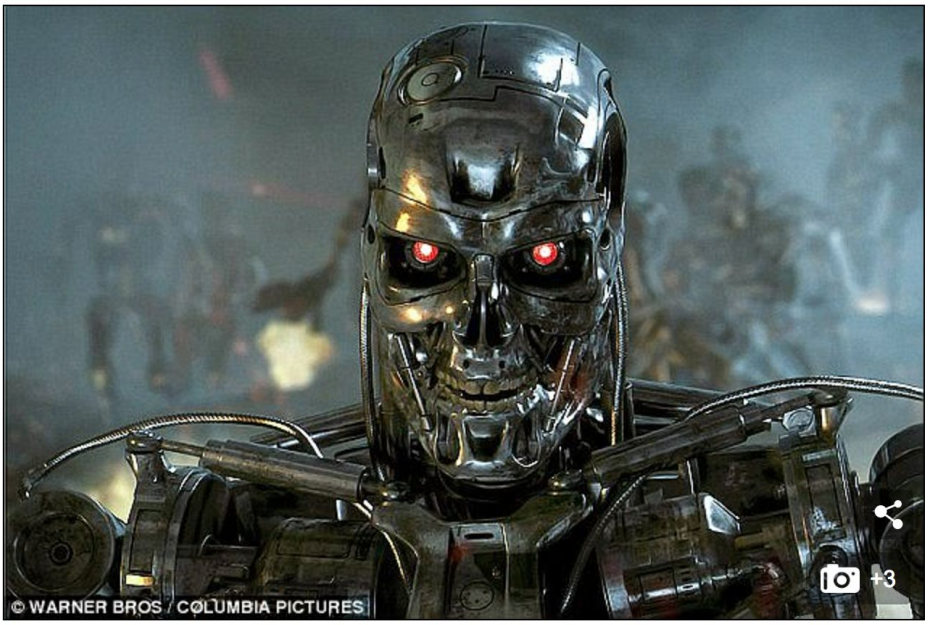

Researchers have long-looked to the human brain for inspiration in creating an intelligent machine that can learn. The machines in the Terminator film franchise (pictured) learn to upgrade and repair themselves and rise up against humanity

This is a actual picture used in an article about AI. If we give AI the ability to learn, logically the next step is that they're going to rise up against humanity!

AI is often depicted as malicious in the media.
Source:
http://www.dailymail.co.uk/sciencetech/article-4382162/Scientists-create-AI-LEARNS-like-human-mind.html

NEWS & TECHNOLOGY 22 February 2017

AI learns to write its own code by stealing from other programs

What do you think of AI after reading this headline?

- AI is malicious. It STEALS.
- Secondly, we LOVE to personify AI. This headline talks about AI like it's sentient, with the intention to learn. One day, AI woke up and decided to teach itself how to code.. And how did it decide to do this? By STEALING from other innocent programs.

Why do we do this? It's somewhat understandable, since headlines which prey on our fear are more interesting, compared to one that says "Grad student improves the accuracy of a program by 2%"

Source:
https://www.newscientist.com/article/mg23331144-500-ai-learns-to-write-its-own-code-by-stealing-from-other-programs/

But should we be alarmed that some of the smartest people on the planet are raising similar fears?

http://observer.com/2015/08/stephen-hawking-elon-musk-and-bill-gates-warn-about-artificial-intelligence/

Elon Musk feels that AI is our greatest existential threat.

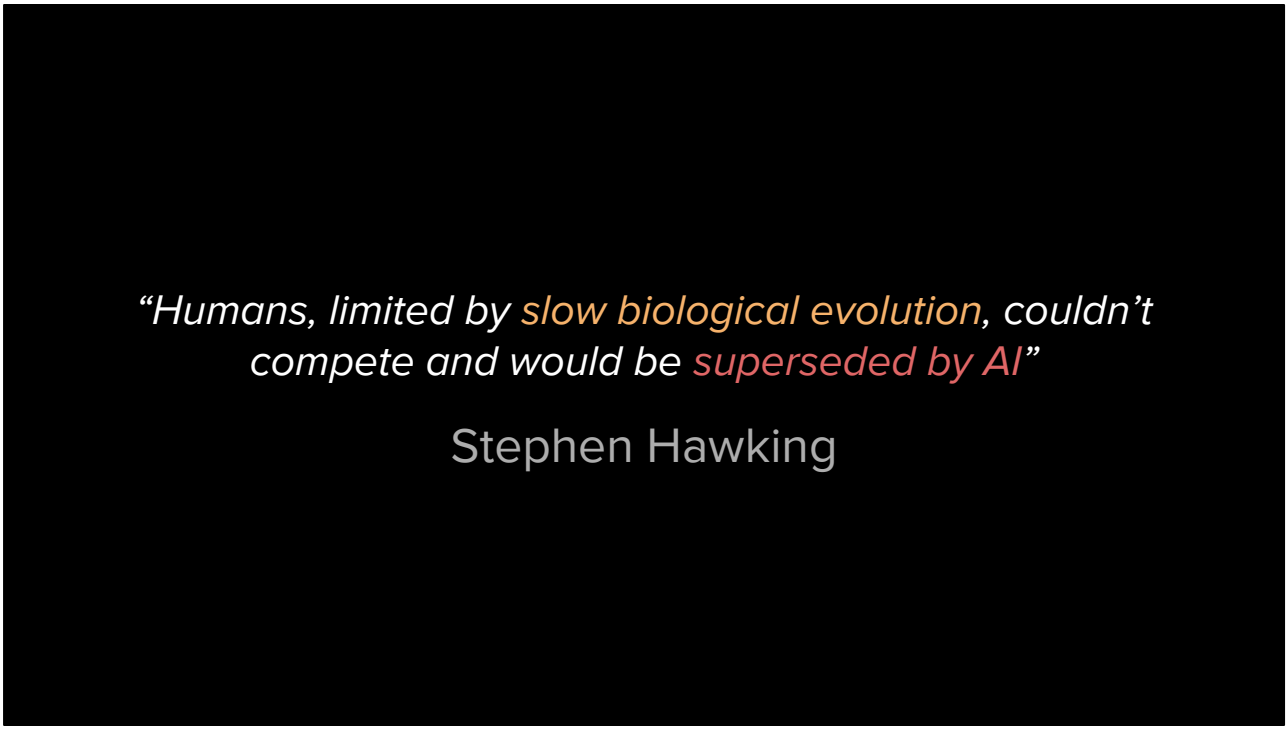

Meanwhile, Stephen Hawking thinks AI will out-evolve humans, and we will be left behind

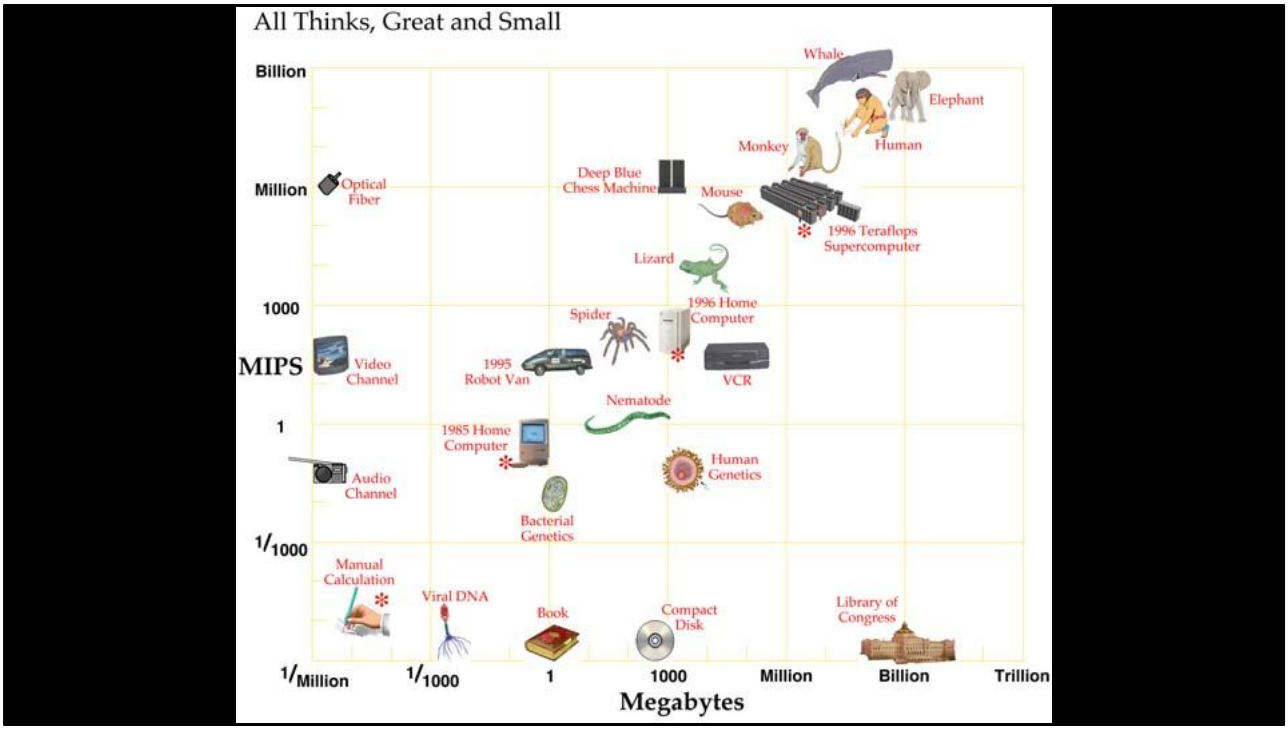

And this may be grounded in truth.
Take a look at this chart - on the y axis, we have processing power - how many instructions can be executed per second. On the x-axis, we have memory capacity. Near the lower left, we have the 1985 home computer, which had the memory and processing power roughly equivalent to a single cell bacteria. Towards the upper right, we have Deep Blue, which can beat humans at chess.
Think about how long it took evolution and natural selection to go from a single celled organism to something as complicated as the mammalian brain. Then think about how it only took three decades to go from windows 1.0 to the kind of computers we have today, which beat us at Go, and can store the entire library of congress.
If we keep up this rate of development, **it's not entirely unfathomable that computers will exceed human intelligence, and achieve a level of complexity of thought that the human mind cannot comprehend.**
http://www.frc.ri.cmu.edu/users/hpm/book97/ch3/index.html

AI in the Real World

So… Is there cause for worry?
To answer that, let's look at the state of AI in the real world

Weak AI

There are basically three types of AI. The first is weak AI, which is non-sentient, and very good at one particular task, and nothing else.
https://en.wikipedia.org/wiki/Weak_AI

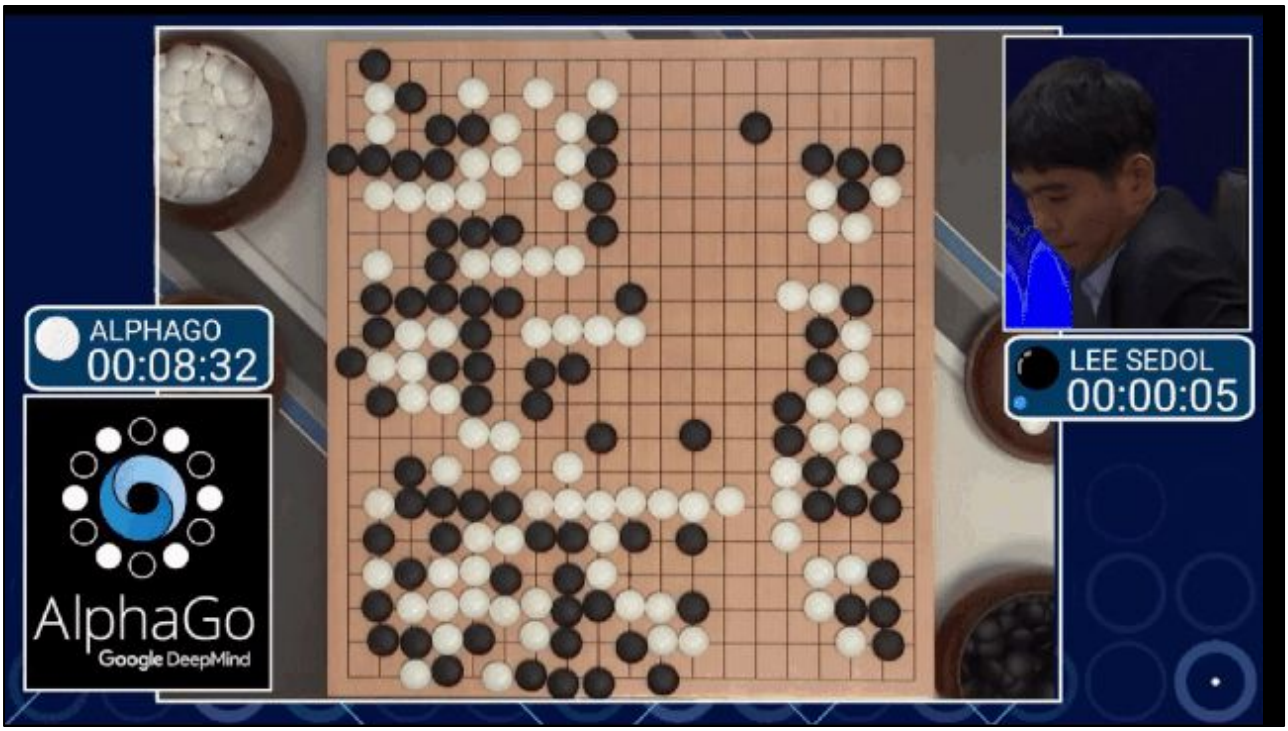

AlphaGo is an example of weak AI - it’s extremely good at one thing, playing Go, but it can’t really do anything else.

Weak AI

Strong AI

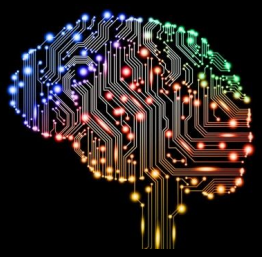

Then there's Strong AI - the kind that is frequently depicted in pop culture, that is, AI that is sentient, with true intelligence, at the level of a human.

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

Weak AI Strong AI

Super Intelligence

There’s a third class: “Super intelligence” - when AI surpasses human intelligence, we literally can’t even comprehend this because our brains are not sophisticated enough

Unexcitingly, the vast majority of AI now is weak AI.
Although Weak AI is called “weak”, it is replacing lots of jobs

Sometimes this is good, for example when we replace sweatshops* with...

Ironically, since human labelers/trainers are required to train algorithms like those used in computer vision, a whole new category of tedious, low wage job has been created, e.g. Mechanical Turk

High tech “sewbots”

https://techcrunch.com/2016/10/09/industrial-robots-will-replace-manufacturing-jobs-and-thats-a-good-thing/
https://thebossmagazine.com/sewbots-are-coming-to-a-closet-near-you/

Another example is truck drivers.
They're overworked, fatigued, and involved in 500,000 truck accidents a year in the US alone,

http://www.truckaccidentattorneysroundtable.com/blog/3-reasons-truck-drivers-continue-to-drive-fatigued/
http://www.chemomentum.org/tips-to-stay-awake-during-a-long-haul/

We can replace these with self driving trucks, that do not get fatigued, do not have to leave behind their families for weeks, eating processed food for months on end, while getting almost zero exercise or mental stimulation..

AI is mostly replacing jobs we don't want, are dangerous, or with poor conditions for humans.

https://www.wired.com/2015/05/worlds-first-self-driving-semi-truck-hits-road/

But not only blue collar jobs are being replaced.
Certain prestigious, white collar jobs are being threatened too, especially quantitative roles like traders and hedge fund managers

http://static1.businessinsider.com/image/51b8de656bb3f7095300000a/why-you-should-try-being-a-trader-at-least-once.jpg

CADE METZ BUSINESS 01.25.16 7:00 AM

THE RISE OF THE ARTIFICIALLY INTELLIGENT HEDGE FUND

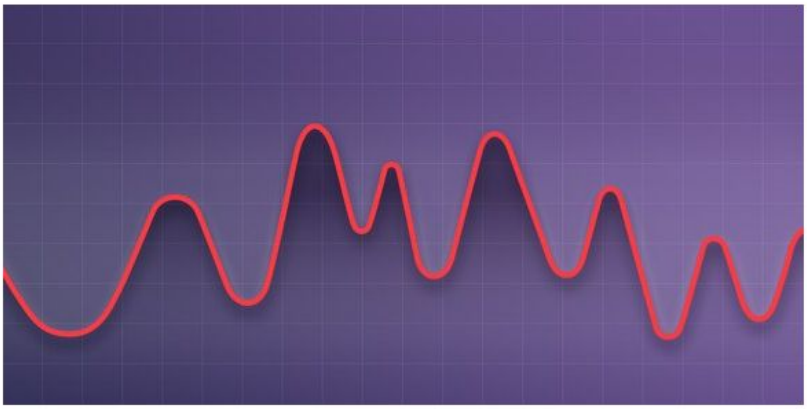

THEN ONE/WIRED

These roles are being replaced by engineers who develop trading algorithms, which are unbiased, unemotional, way faster and better at numbers

https://www.wired.com/2016/01/the-rise-of-the-artificially-intelligent-hedge-fund/

AI + Humans

So with all these jobs being replaced, is there any room for humans left?
Despite the bleak outlook, there are actually many problems which still require human intelligence
In most situations, we'll most likely use AI as a tool to help us do our existing jobs better

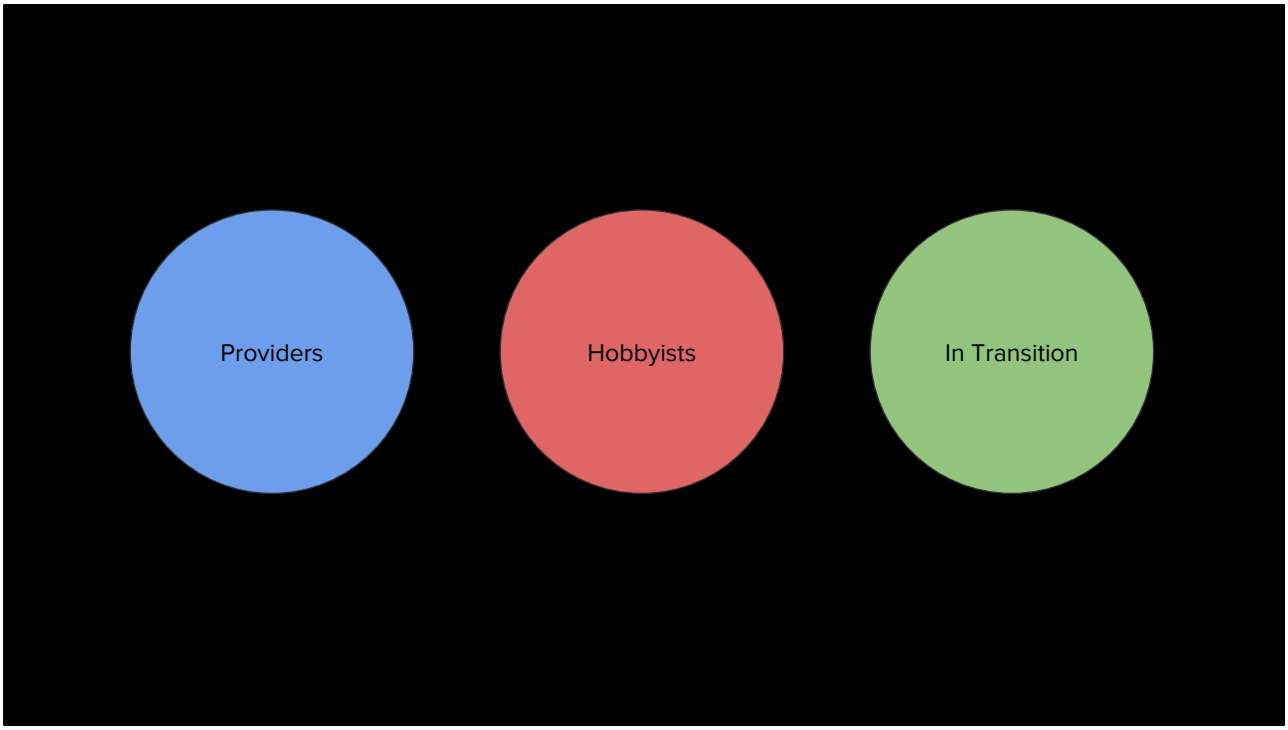

You'll probably be familiar with many of these examples, such as using AI to segment our customer base.
Clustering algorithms can be used to consider hundreds of dimensions simultaneously to cluster groups of people together.
For example at Uber, we use this to cluster our driver partners into distinct groups, each with distinct behaviors and needs, so that we can better serve each unique segment

We also have algorithms that can help us assess the effectiveness of content.
When we come up with a ad - how do we tell its good? We can invite focus groups, or we can click a button and use algorithms which can predict which parts of an image people tend to look at - you don’t even need real people to measure this anymore.

We can also detect emotions from videos. Humans can train for years in order to pick up on microexpressions that last a fraction of a second, and may sometimes be below our conscious awareness. Imagine what a computer capturing a million frames per second would be able to do. They could potentially pick up on emotions even we aren't aware we're experiencing

We're just getting started in this field and it's already pretty good.

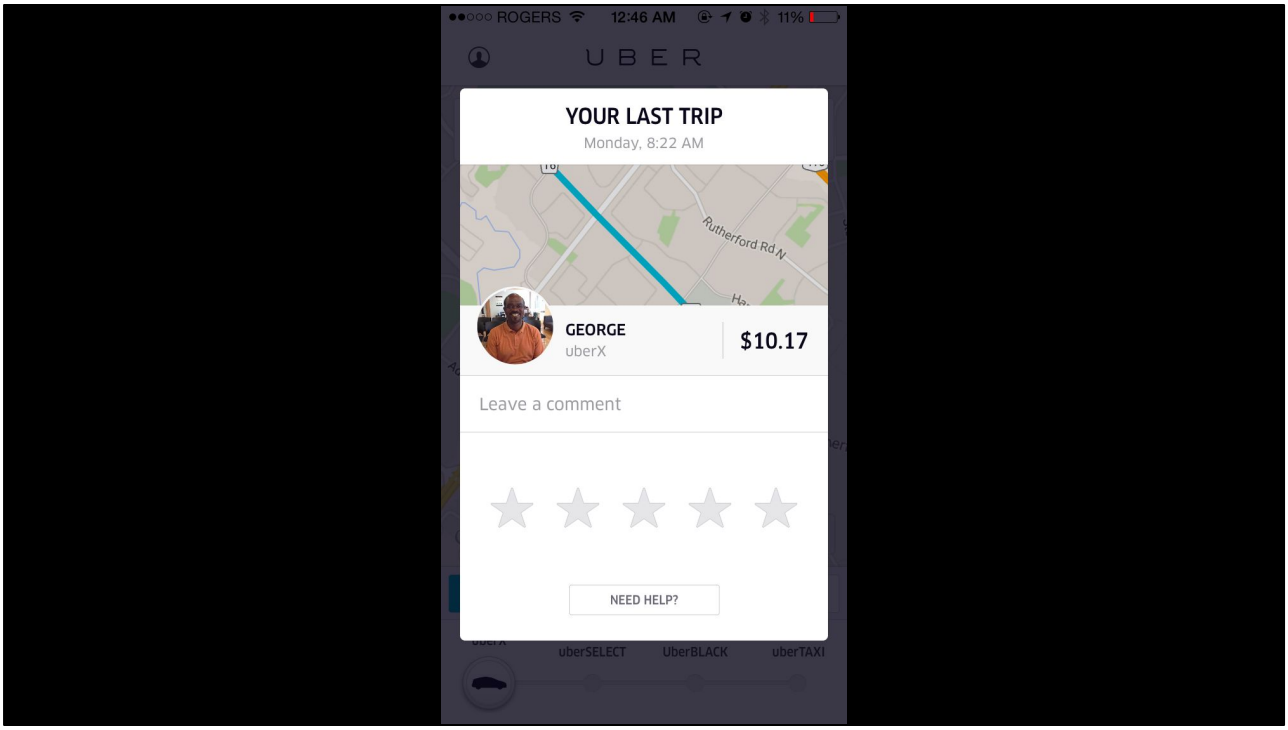

We can also extract meaning and sentiment from language
(pic source: http://bramptonist.com/uber-taxo-bramptn/)

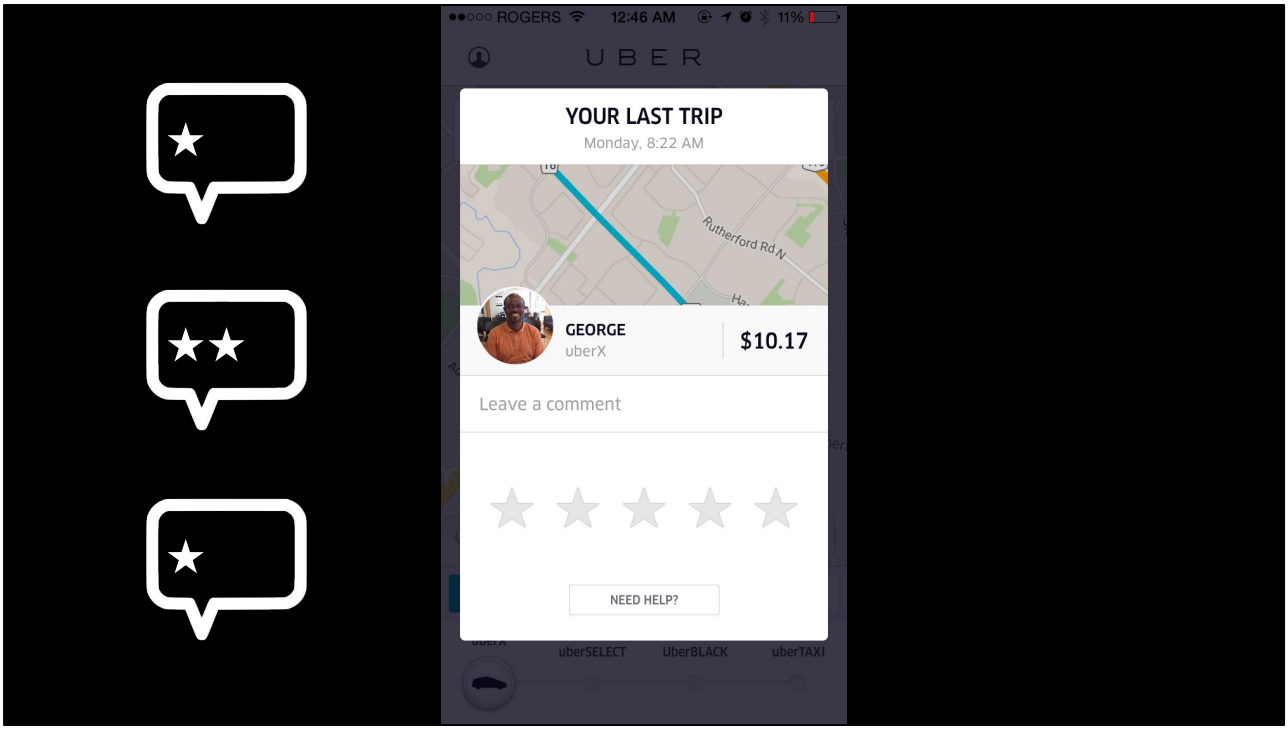

Say we get text feedback from trips

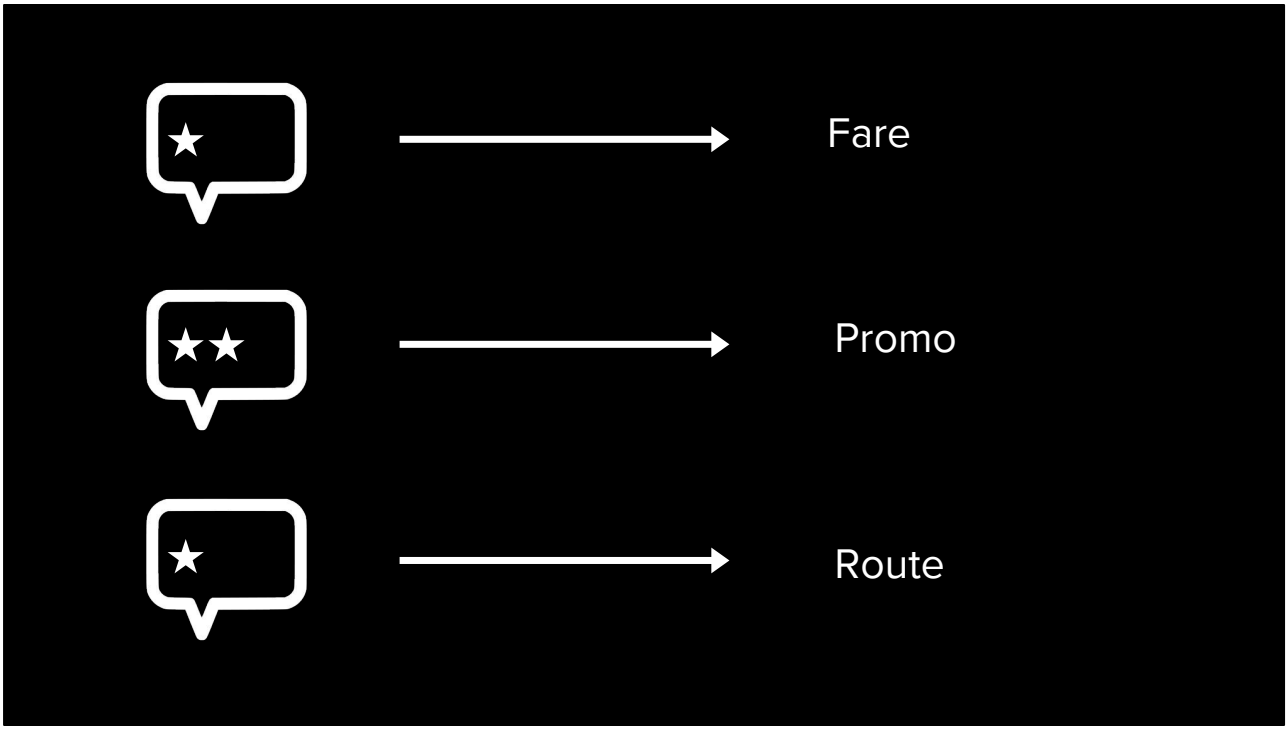

Natural Language Processing, as the name suggests, is a branch of AI which deals with language

It can be used to extract keywords and classify the issues, and also derive the sentiment.

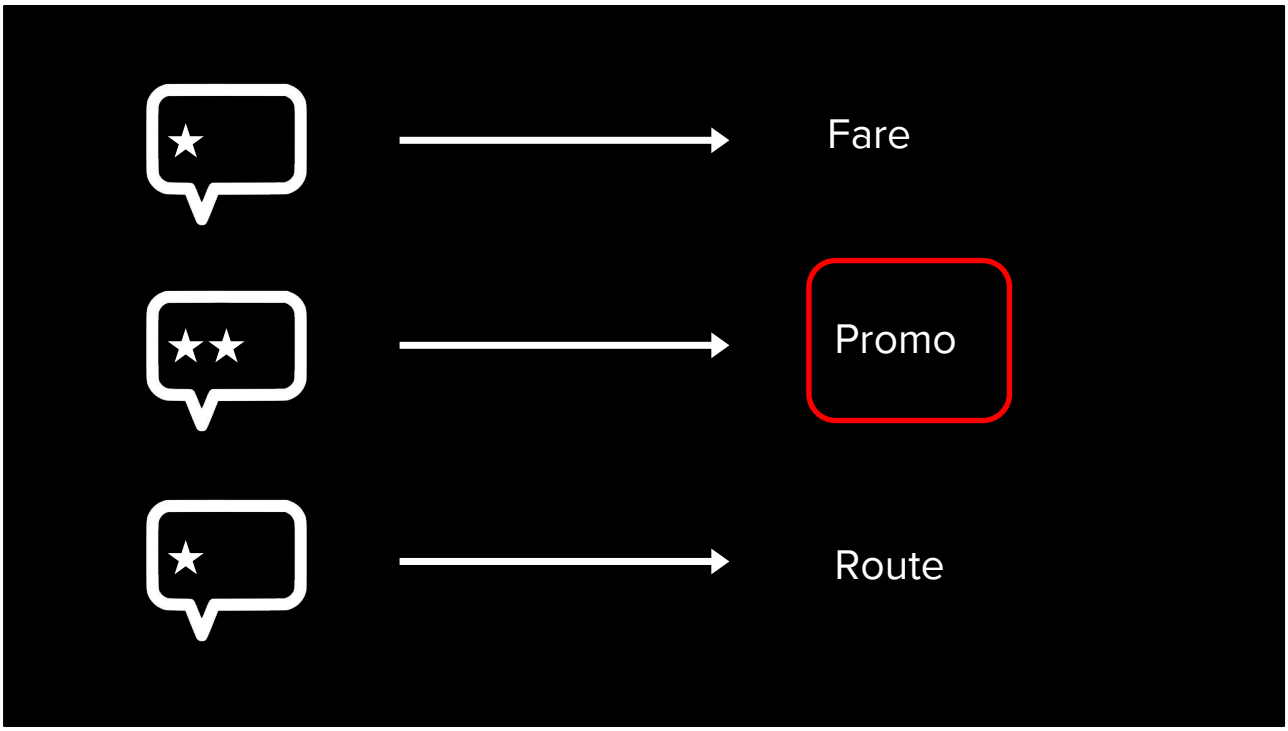

We can then more easily extract insights from a large number of comments.
E.g. if I'm on the team that cares about promos, we extract all the promo comments, then perform sentiment analysis to find out if the general sentiment is positive or negative, without having to make humans go through each comment.

AI’s limitations

There are endless examples I could raise, but with this small number of examples, we can already see that the potential of AI as a tool is huge, helping us generates insights not previously possible.
However, despite AI’s power and potential, it comes with many limitations as well

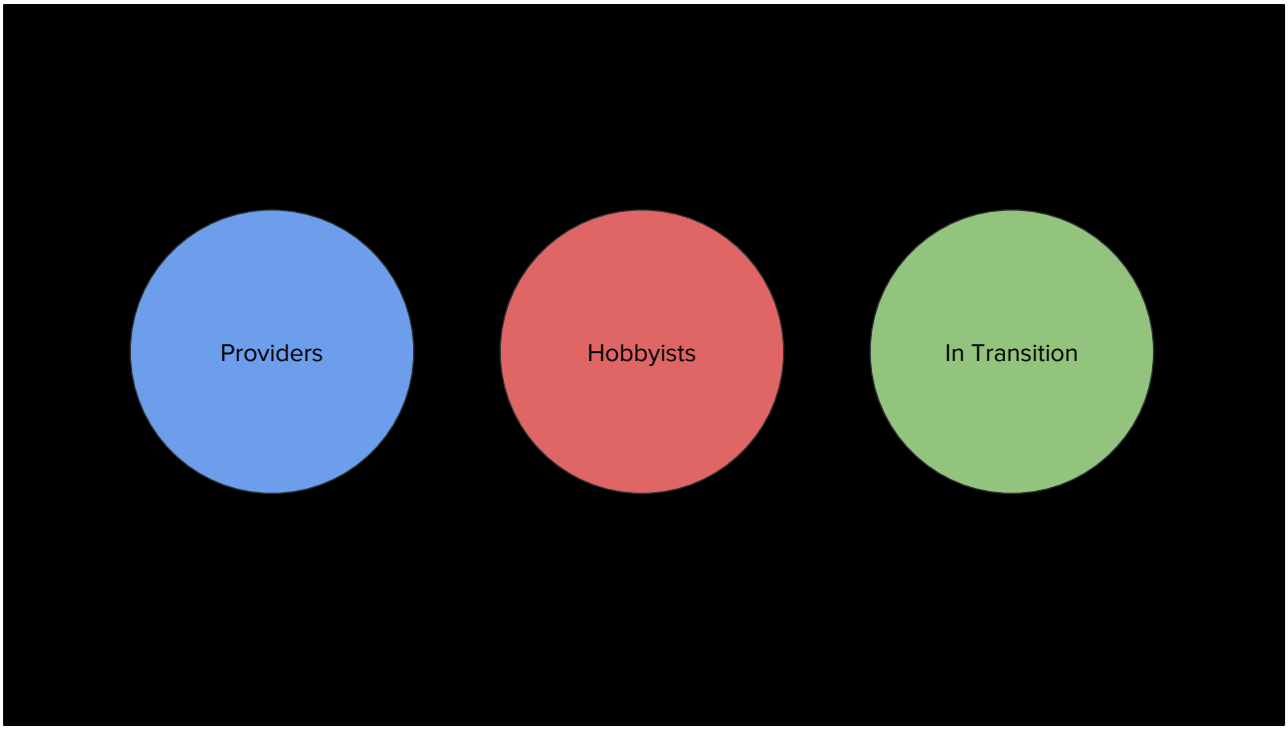

Just now I provided this example of clustering algorithms identifying segments for us
This is what we usually see when teams present their findings to us - clusters come with nice labels and interpretations

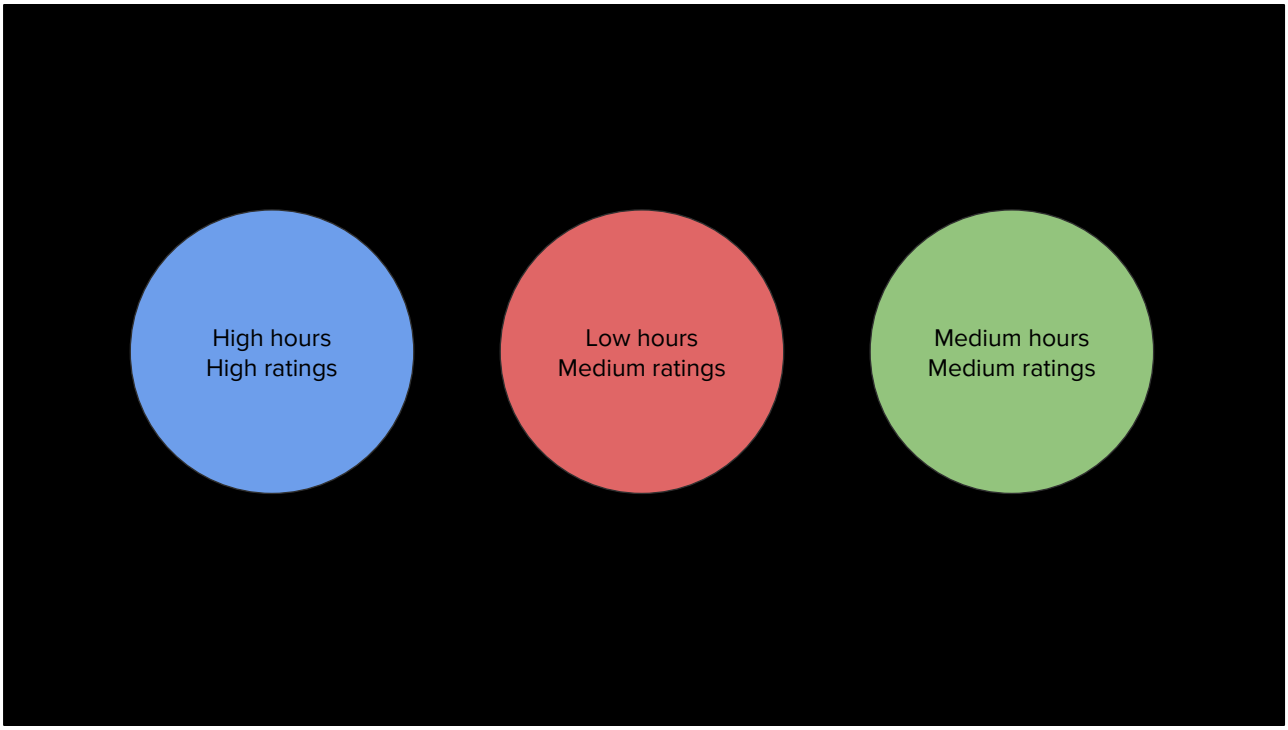

In reality, raw results from the algorithms that look more like this, just a bunch of numbers.
It's all pretty meaningless until a domain expert interprets the results and assigns meaning to the raw numbers.
This is just one example, but in many different types of analyses, AI is great at the quantitative side of things, but fails at the qualitative aspects of the task, such as generating labels.

Language generation is another qualitative task that AI fails at quite miserably.
Earlier I brought up the example of NLP.
While AI is pretty good at parsing language, it's not very good at generating language.
When it comes to even simple things like coming up with names for paint colors, it's pretty bad - look at the results of a neural network that was trained to generate color names for paint.

https://arstechnica.com/information-technology/2017/05/an-ai-invented-a-bunch-of-new-paint-colors-that-are-hilariously-wrong/

```
INT. SHIP

We see H pull a book from a shelf, flip through it while
speaking, and then put it back.

                    H
          In a future with mass unemployment,
          young people are forced to sell
          blood. That's the first thing I can
          do.

                    H2
          You should see the boys and shut
          up. I was the one who was going to
          be a hundred years old.

                    H
          I saw him again. The way you were
          sent to me... that was a big honest
          idea. I am not a bright light.

                    C
          Well, I have to go to the skull.  I
          don't know.

He picks up a light screen and fights the security force of
the particles of a transmission on his face.
```

video

It gets even worse if we try to generate longer, more complex content. Some people tried to train a neural network to write a screenplay, and this is the result.

AI is able to string words together according to grammatical rules, but not in a way that makes sense.

This is an excerpt from “Sunspring”, a screenplay that was written by a neural network (this was trained on human screenplays, or as we call them, screenplays.)

Sunspring - written by Benjamin, a “self-improving LSTM RNN machine intelligence trained on human screenplays”

https://www.youtube.com/watch?v=LY7x2Ihqjmc
https://www.docdroid.net/lCZ2fPA/sunspring-final.pdf.html#page=2

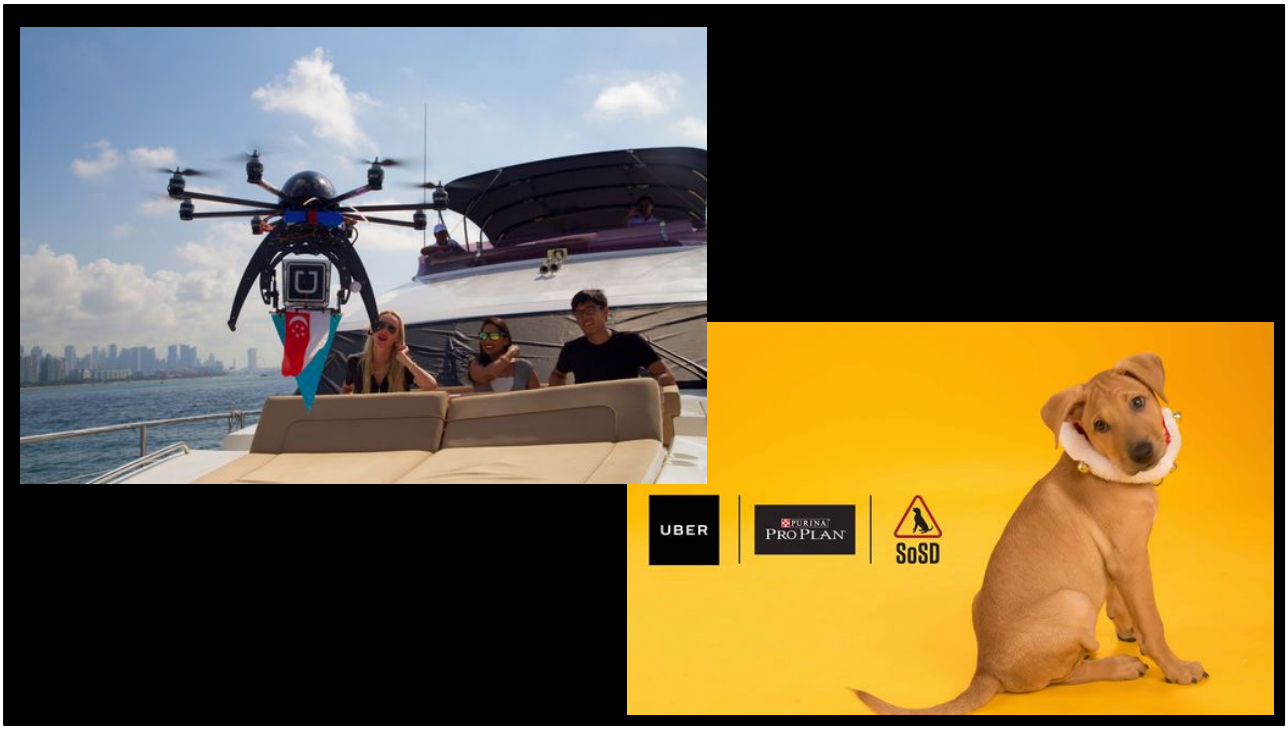

And if we move on from text-based screenplays to even more complex concepts, such as coming up with marketing campaigns like drone delivered ice cream or puppies for christmas, **AI is still very far off from being able to do that**.
Most algorithms are trained to come up with new permutations of things it has seen in the past, but it still takes a human to be truly innovative.

Human Computer Interaction

Speaking of creativity, AI actually creates entirely new design challenges and opportunities for us to be creative.
Because of the atmosphere of fear that AI poses, and also because many AI centered technologies and completely new things that we've never seen before,
There are many design issues to consider, and this has resulted in the explosion of interest in the field of HCI.

For example, say we're designing a self driving car.
I suspect that most of us in this audience would think this car is cool and futuristic, and given the chance to try it out, we wouldn't even hesitate.

But there are still segments who are fearful of new technology, and for them, we need to design to make tech appear less intimidating.
Waymo took an interesting approach by designing their cars to appeal to the human evolutionary tendency to like cute, juvenile features. For example the car has roundish eyes, smaller snout, and it's generally more squished, sort of like the features you'd see in a young mammal.

http://fortune.com/2016/10/12/google-yoojung-ahn-car-designer/
https://www.theverge.com/2017/4/25/15415840/waymo-self-driving-minivan-early-rider-phoenix

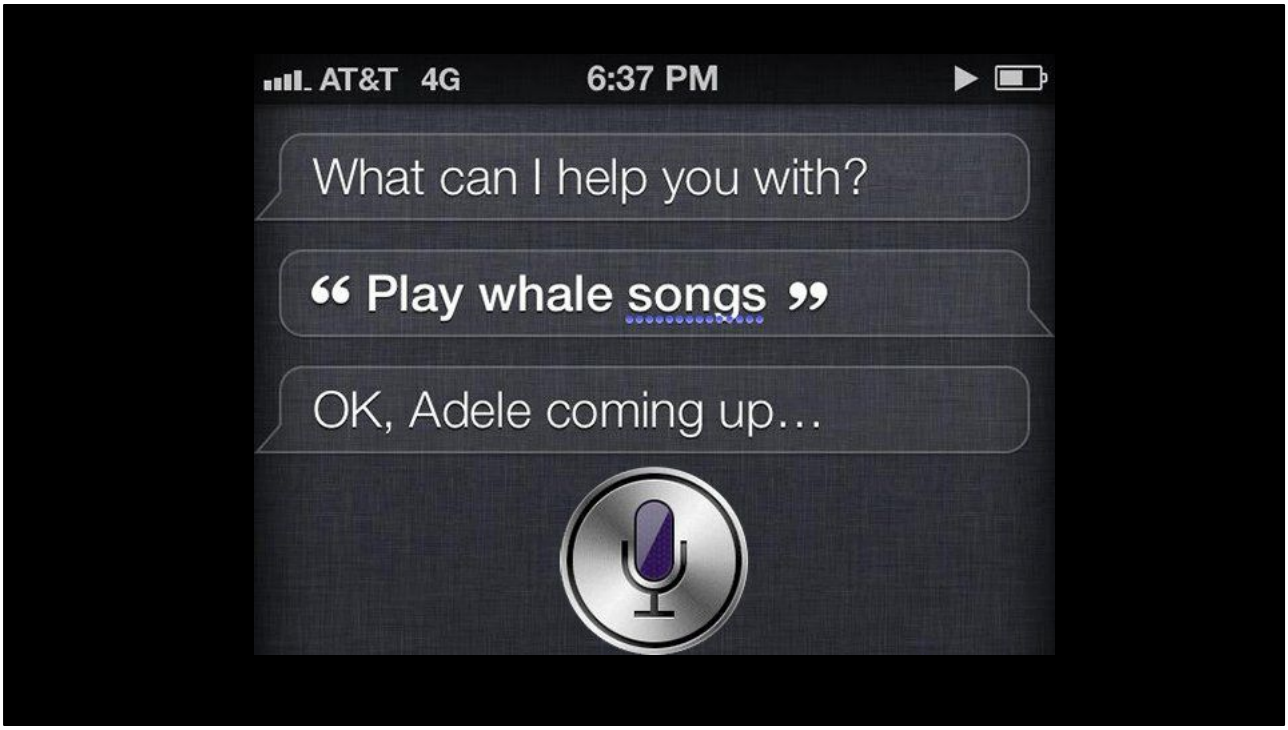

Another example of HCI at work are intelligent personal assistants like Siri.
Apple has done a wonderful job of making Siri seem more like a real person than a robot, by doing things like giving it a sense of humor.
User experience, user-centered design, and human computer interaction are going to be super hot fields as more AI- powered technology appears in the next few decades.

To summarize, AI takes away some jobs, but creates new opportunities.

AI...

Good at:

- Solving very specific, well-structured, quantitative problems

AI...

Bad at:

- Qualitative tasks
- Labeling / Interpretation
- Content creation
- Human-friendliness

Humans

“Piloting” the AI:

- Identifying the problems to solve,
- Knowledge of algorithms,
- Interpreting results

Considering psychological, moral, and emotional aspects

Creativity

Empathy

AI... best used in conjunction with human intelligence

In conclusion, in most cases, AI is best used in conjunction with a human brain to steer it.
Because of AI's limitations and the new problems AI introduces, humans are not yet completely dispensible…

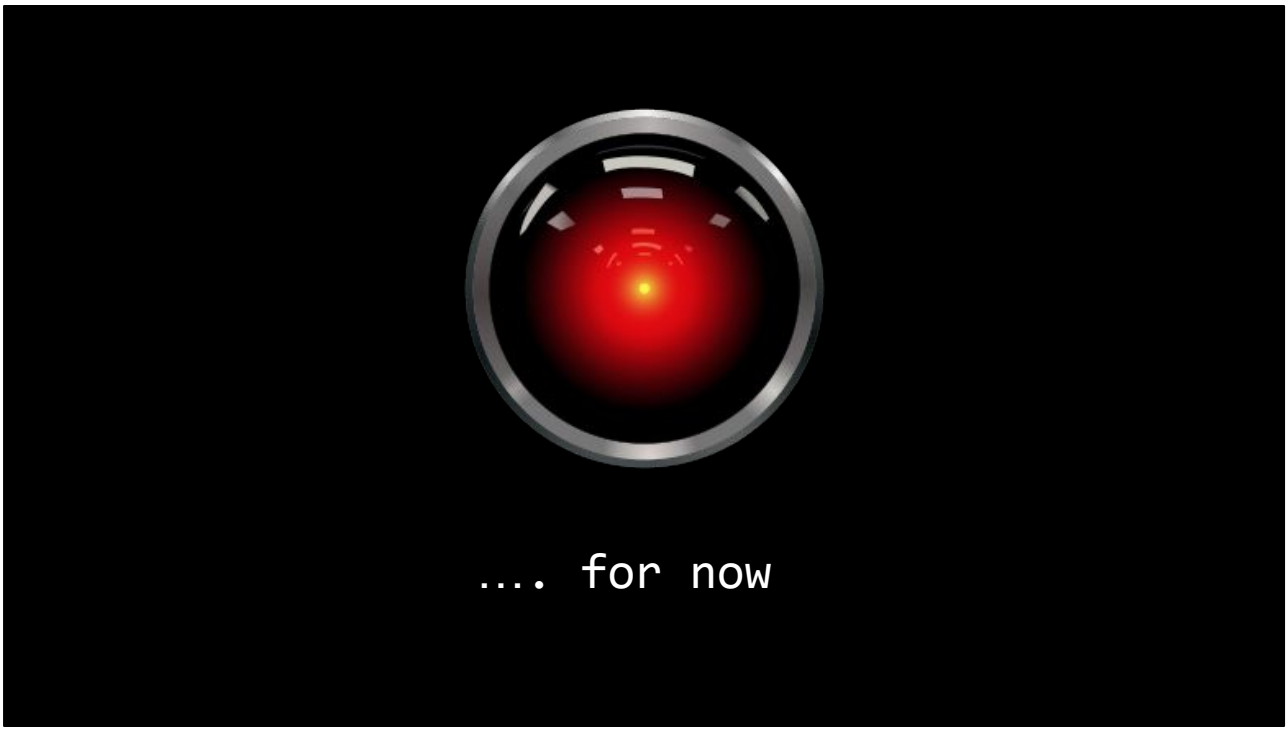
.... for now

mirri66.github.io
www.linkedin.com/in/tsmgrace
grace.tang@uber.com

Grace Tang, Data Scientist
UBER