CS 343

AI: Ethics and Society

Prof. Yuke Zhu

The University of Texas at Austin

Logistics

- Writing responses - Chapter 3 (Next Monday 1/24, 3:00pm)
- Homework 1: Search (Monday 2/7, 11:59am)
- Programming Assignment 1: Released now (Wed 2/9)

PEAS Description of the Task Environment

- An **agent** is an entity that perceives and acts.
- A **rational agent** selects actions that maximize its (expected) **utility**.
- Characteristics of the **Performance measure, Environment, Actuators, and Sensors (PEAS)** dictate techniques for selecting rational actions

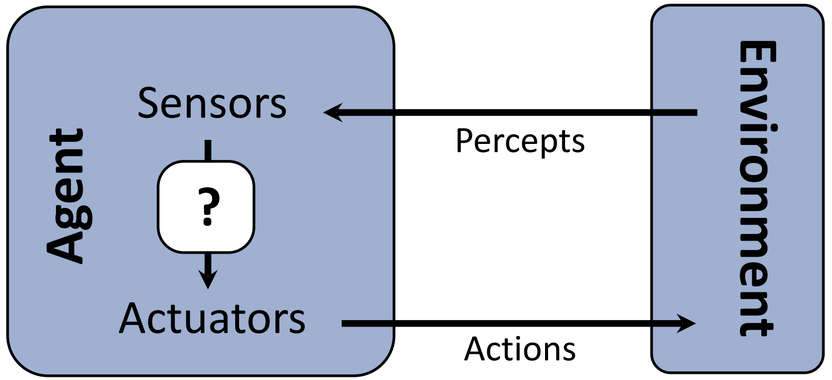

PEAS Description of the Task Environment

- An **agent** is an entity that perceives and acts.
- A **rational agent** selects actions that maximize its (expected) **utility**.
- Characteristics of the **Performance measure, Environment, Actuators, and Sensors (PEAS)** dictate techniques for selecting rational actions

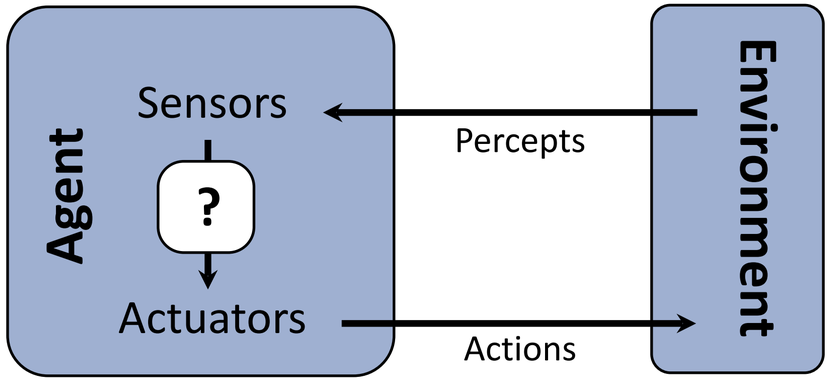

- **(Tyler)** Humanoid robot boxer - **P:** wins, knockouts, and amount of damage, **E:** boxing ring with opponent, **A:** arms, legs, torso, and neck movements, **S:** head-mounted camera and pressure sensors on the body
- **(Mrityunjay, David)** Robotic chef - **P:** Quality of the food and time taken to cook, **E:** Stove and ingredients, **A:** Arms, hands, and legs, **S:** Camera, infrared sensor, joint angle sensors
- **(Risa)** Spam filter - **P:** Amount of spam left in inbox, non-spam falsely marked, **E:** Email inbox, **A:** mark emails as spam, **S:** content in incoming emails

Pac-Man as an Agent

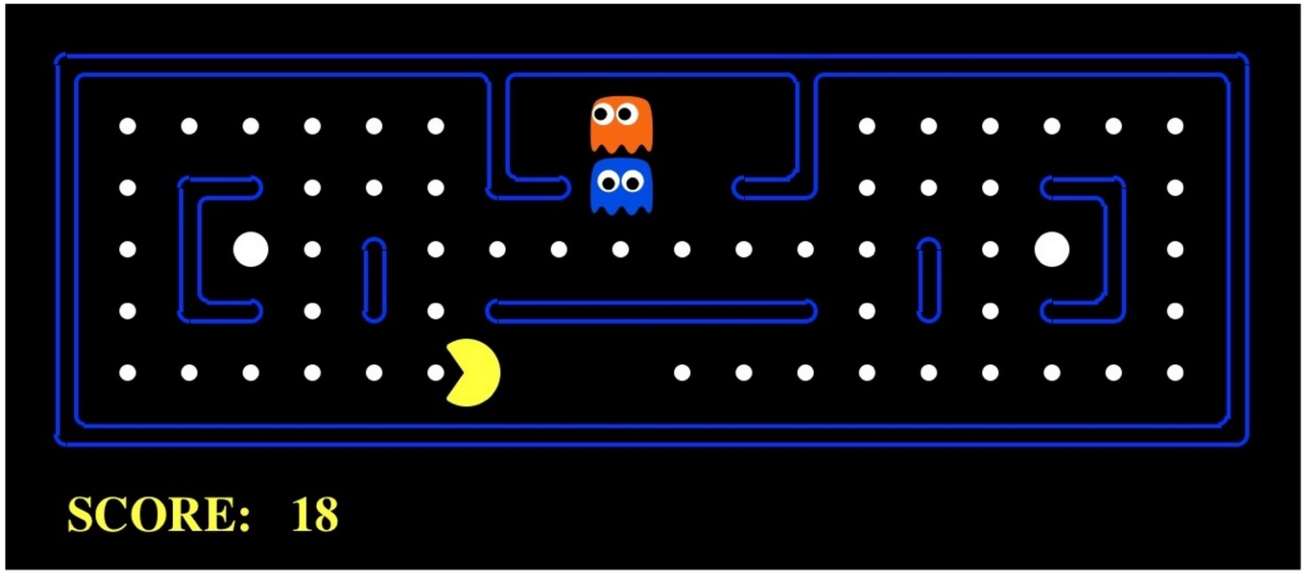

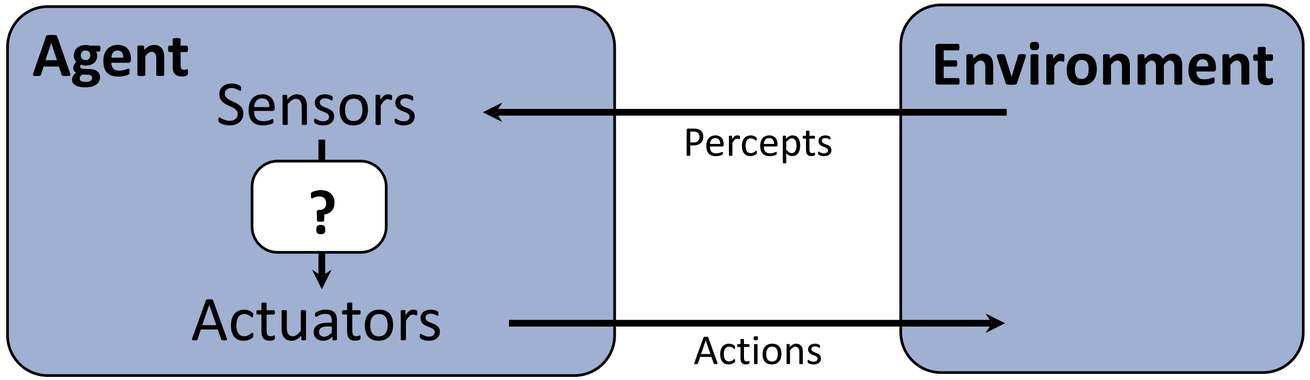

AI Everywhere

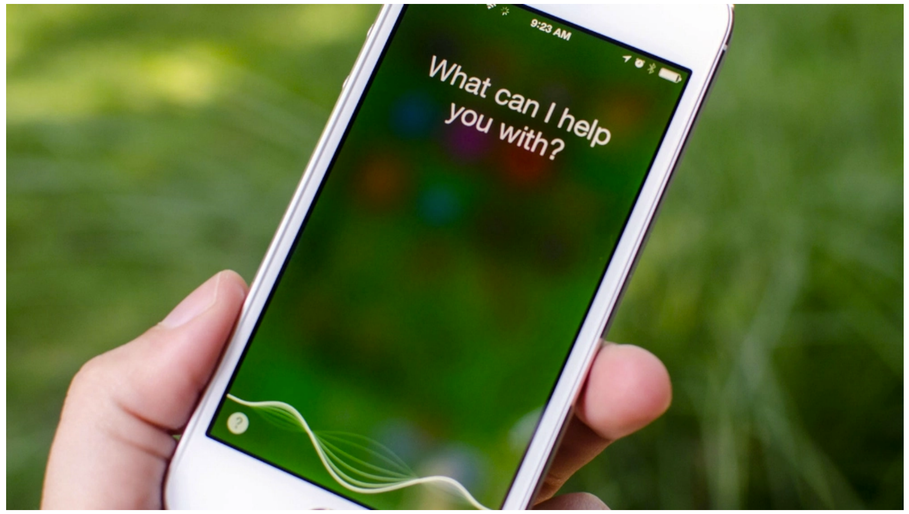

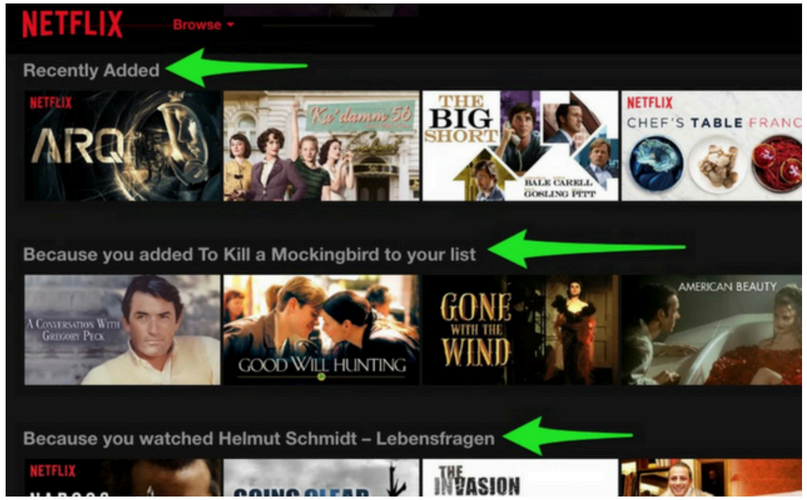

Utopia or Dystopia?

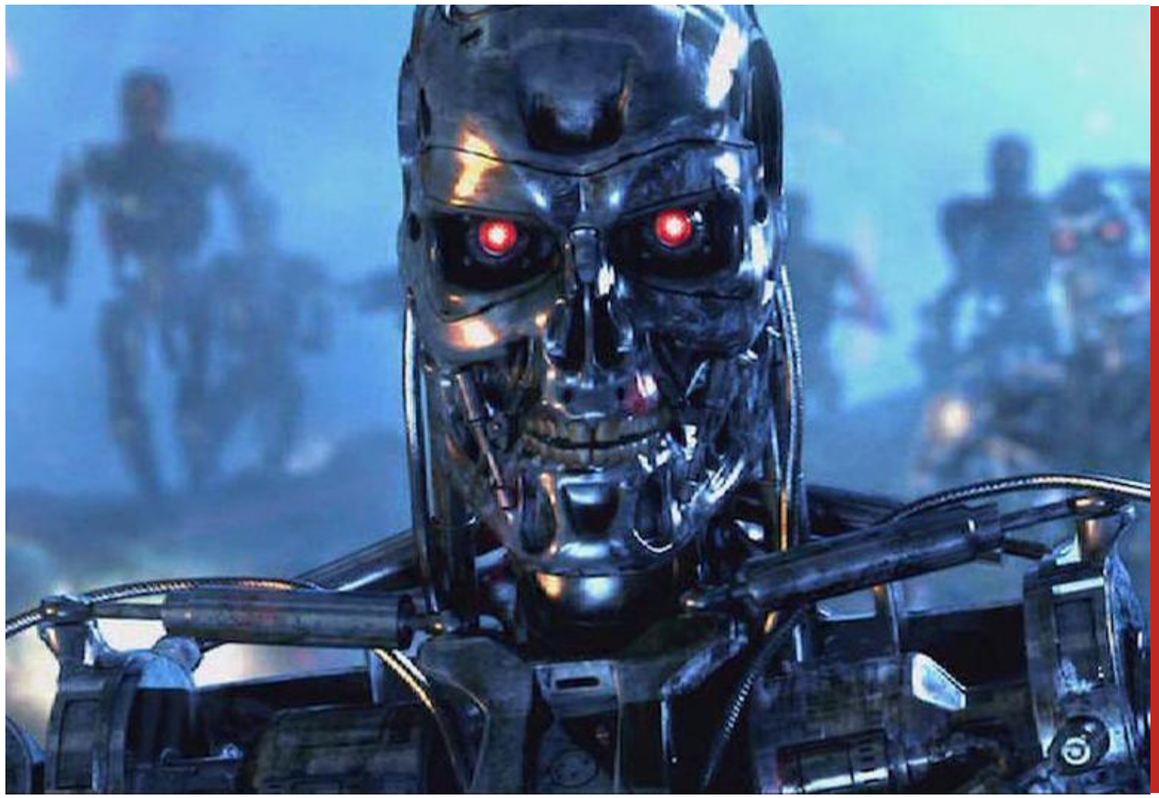

Not so fast…

AI and COVID-19

THE ROBOTREPORT
EXPLORING THE BUSINESS AND APPLICATIONS OF ROBOTICS

RESEARCH TECHNOLOGIES DEVELOPMENT ROBOTS MARKETS INVESTMENTS RESEARCH

Will COVID-19 accelerate an automated future?

By Bastiane Huang | March 29, 2020

BBC Sign in News Sport Reel Worklife Travel Future Mor

NEWS

Home US Election Coronavirus Video World US & Canada UK Business Tech

Technology

Coronavirus: Will Covid-19 speed up the use of robots to replace human workers?

By Zoe Thomas
Technology reporter

19 April 2020 Share

Coronavirus pandemic

WIRED BACKCHANNEL BUSINESS CULTURE GEAR IDEAS SCIENCE MORE SIGN IN SUBSCRIBE

ILLUSTRATION: BETH HOLZER

KAI-FU LEE BACKCHANNEL 05.22.2020 07:00 AM

Covid-19 Will Accelerate the AI Health Care Revolution

Disease diagnosis, drug discovery, robot delivery—artificial intelligence is already powering change in the pandemic's wake. That's only the beginning.

MIT Technology Review

Sign in Subscribe

Topics Magazine Newsletters Events

Artificial intelligence / Robots

Covid-19 could accelerate the robot takeover of human jobs

Machines were supposed to take over tasks too dangerous for humans. Now humans are the danger, and robots might be the solution.

Discussion Question

When will AI reach human-level intelligence?

- In the next 10 years
- In the next 50 years
- In the next 100 years
- Later or Never

Discussion Question

When will AI reach human-level intelligence?

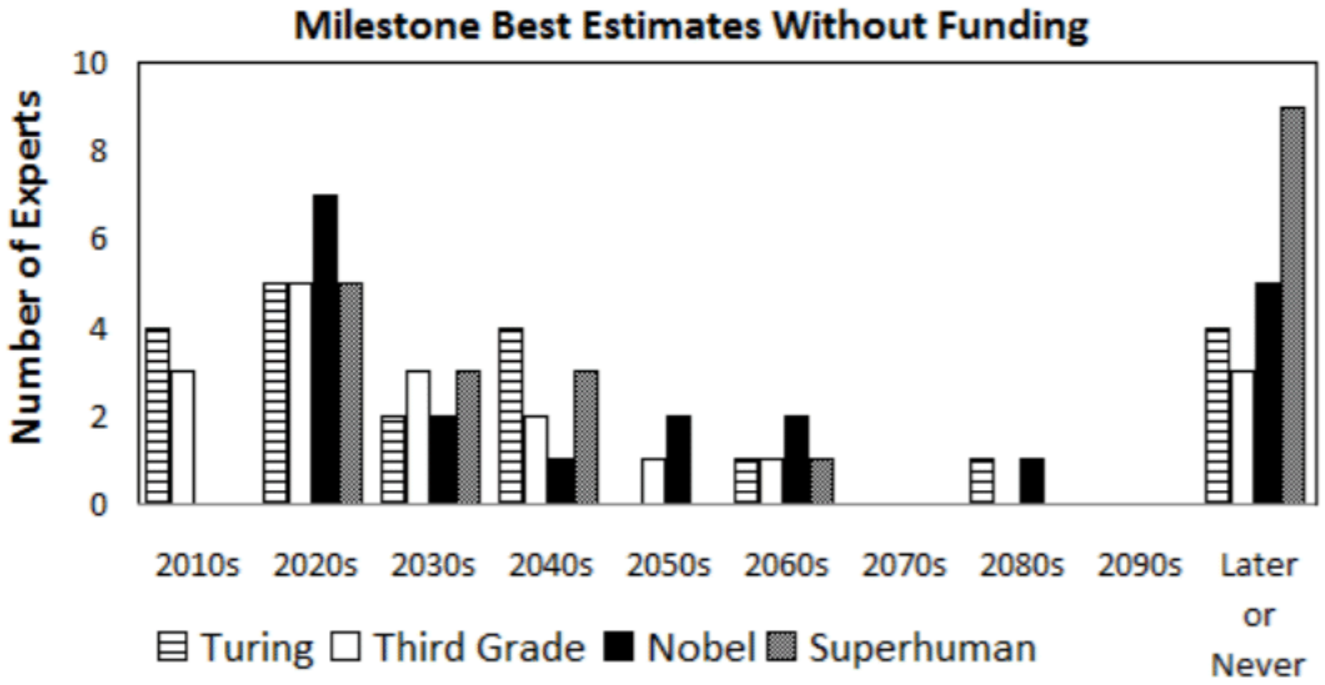

[Source: How Long Until Human-Level AI? Results from an Expert Assessment]

Human vs. AI Characteristics

Human	AI
Evolved for survival	Designed by engineers
Sets own goals	Goals programmed explicitly (usually)
Complex, general purpose system	Specific, constrained system
Continually learns	Can turn off learning, or not use learning
Learns from all observed data	Data access can be controlled
Learns only from own experiences	Can share data with other robots
Can make any choice at any time	Available actions can be restricted

Human vs. AI Characteristics

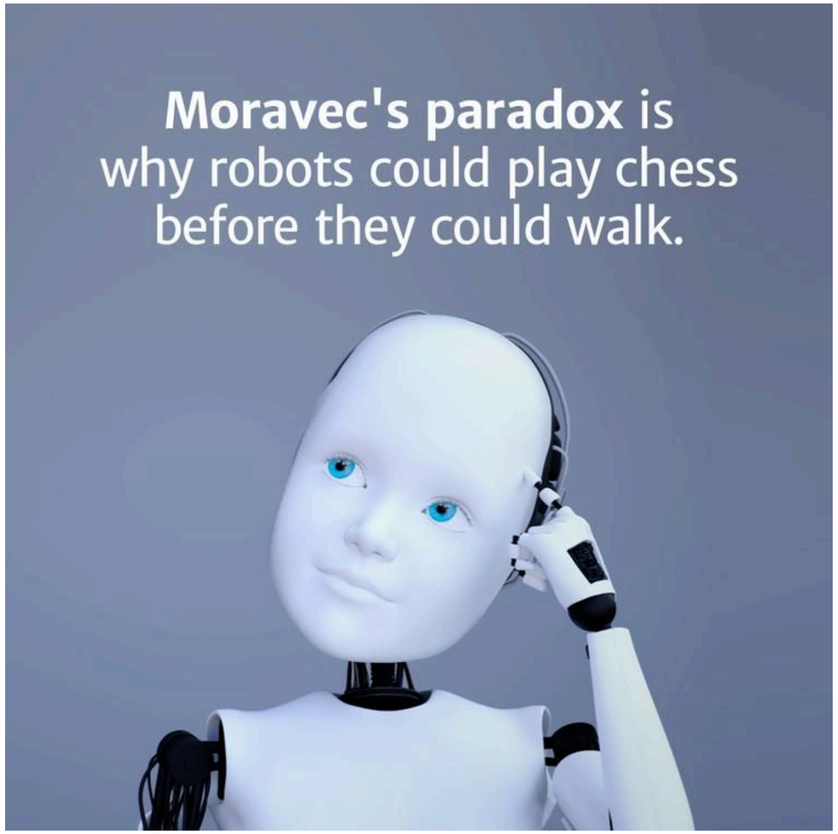

Acting Rationally

Acting rationally simply means maximizing utility

…but can this go wrong?

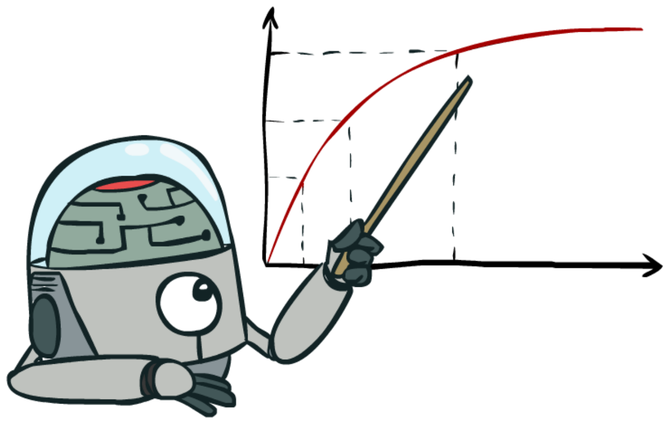

Unforeseen Consequences of Maximizing Utility?

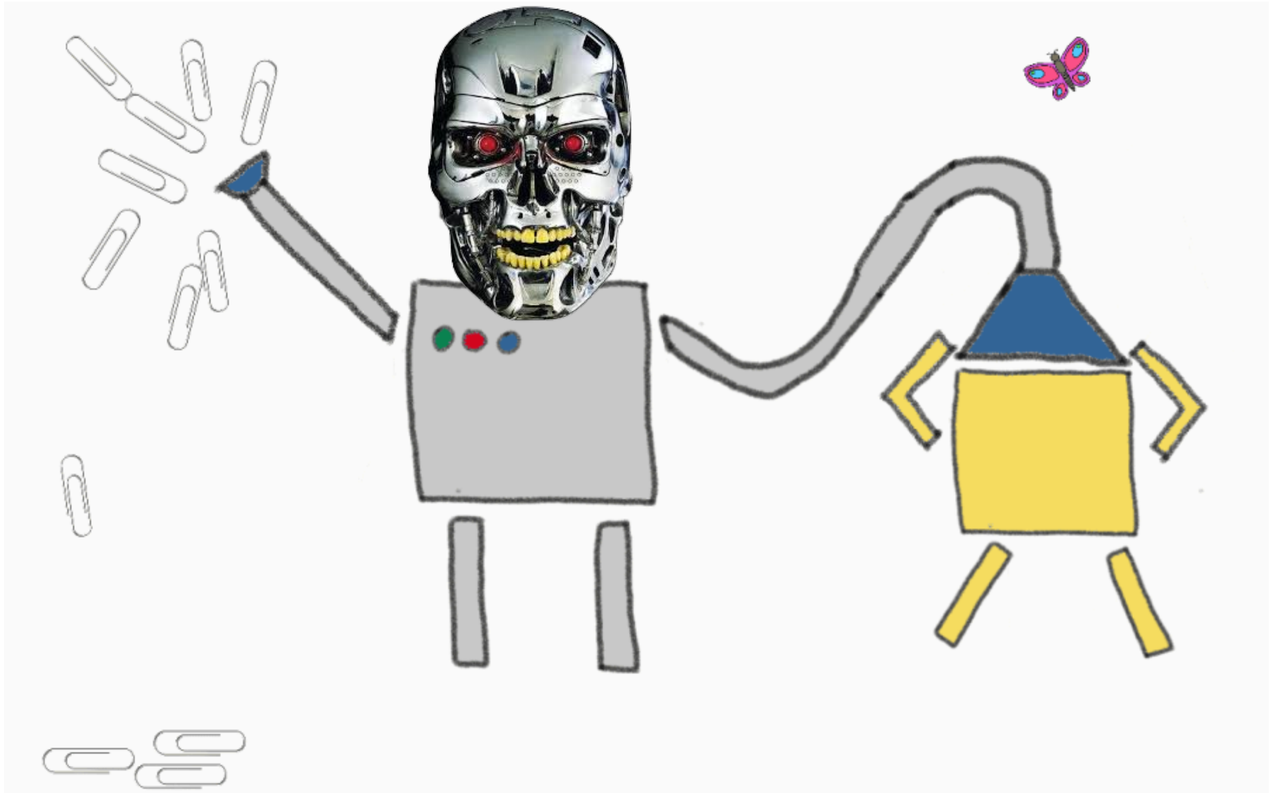

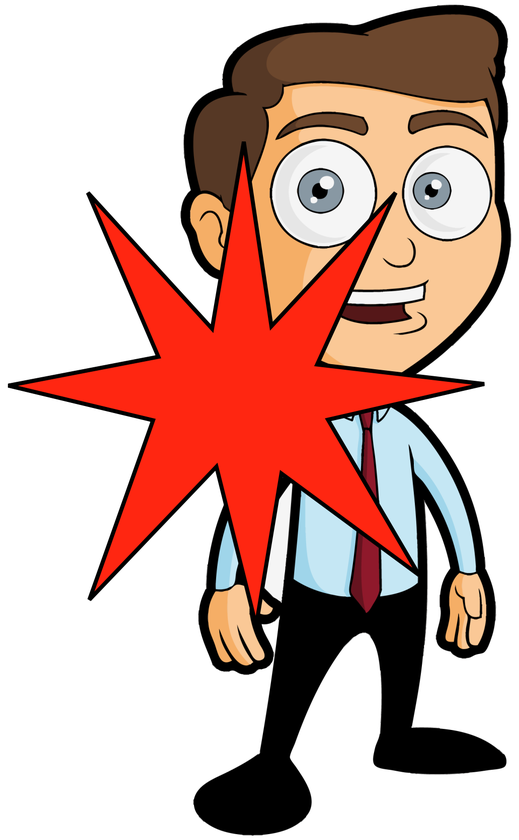

“paperclip maximizer” thought experiment

What went wrong?

- **Is this realistic?**
  - Robots aren't smart enough to be self-aware of their on/off states or to understand chemistry. But let's assume they will be able to in the future.
  - It wouldn't have a concept of "human" to go seek out. It only knows about making paperclips.

- **Bad design!**
  - Objectives must be designed carefully: robot should only be rewarded for making paperclips.
  - Actions should be limited: only actions available should be to make paperclips.
  - Plans should be verified for safety before / during execution: cancel any trajectory that will come in contact with a human.
  - Don't continue learning after deployment.

- **Is this any more dangerous than any factory with non-intelligent machinery that doesn't automatically stop if someone is in the way?**
  - It is bad design, but we know how to use engineering to avoid these situations!

Realistic Risks of AI

Mass unemployment due to automation

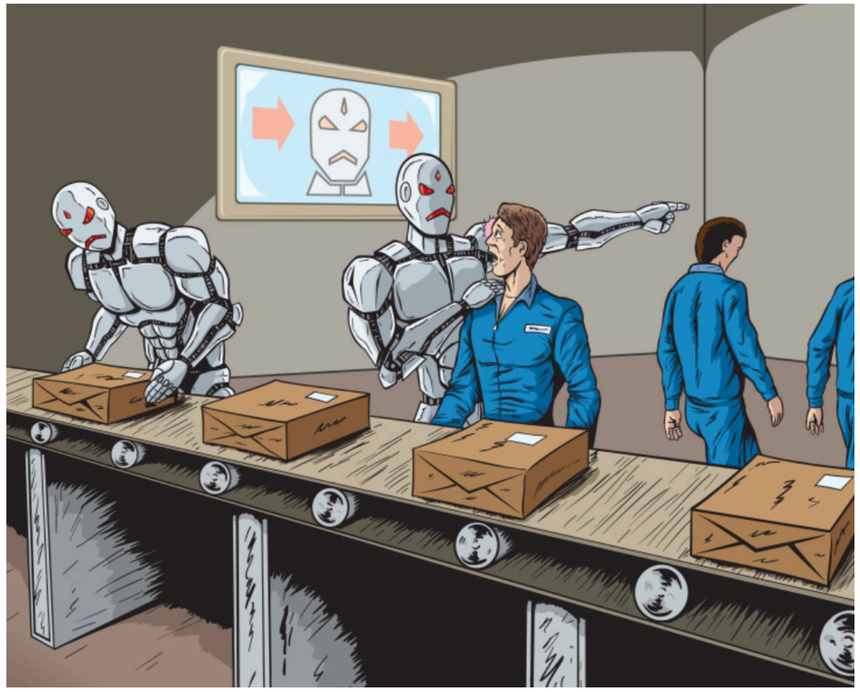

Realistic Risks of AI

Substandard testing / poor user understanding

Realistic Risks of AI

How to make tough decisions?

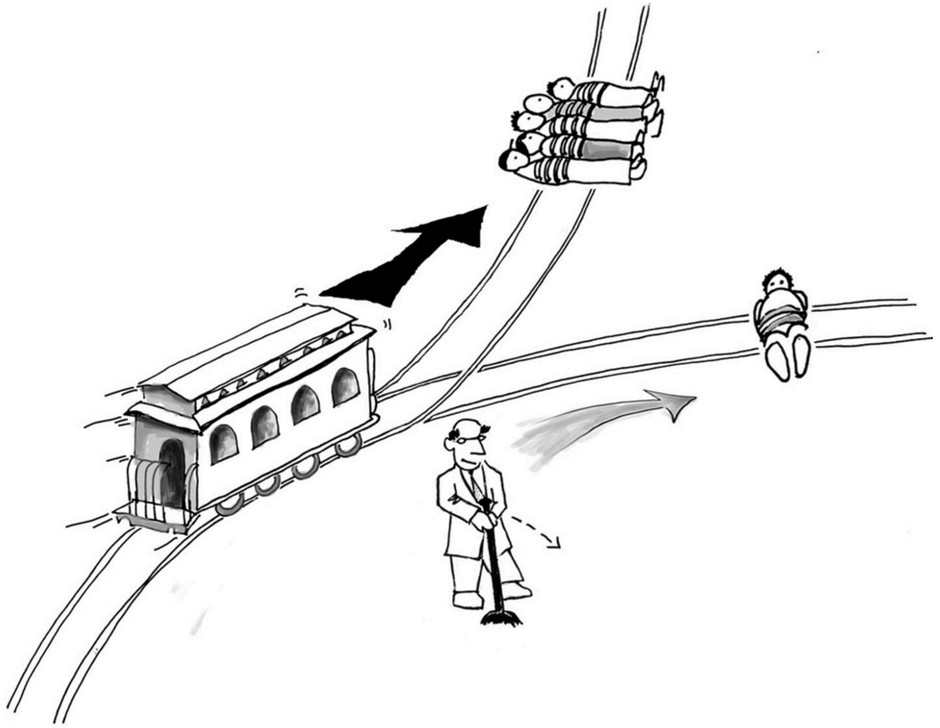

The trolley problem

Realistic Risks of AI

Privacy concerns

The New York Times

Facial Recognition Tech Is Growing Stronger, Thanks to Your Face

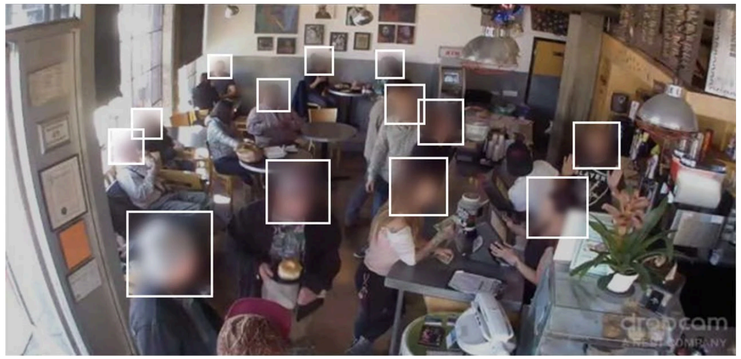

The Brainwash database, created by Stanford University researchers, contained more than 10,000 images and nearly 82,000 annotated heads.
Open Data Commons Public Domain Dedication and License, via Megapixels

By Cade Metz

July 13, 2019

251

Realistic Risks of AI

Algorithmic bias and discrimination

Darker Male	Darker Female	Lighter Male	Lighter Female
94.0%	79.2%	100%	98.3%
99.3%	65.5%	99.2%	94.0%
88.0%	65.3%	99.7%	92.9%

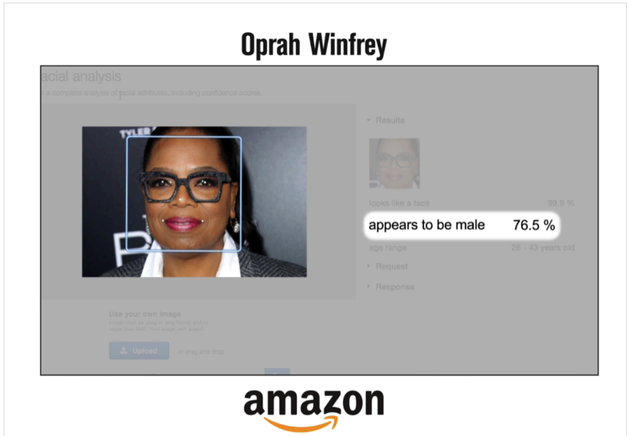

Realistic Risks of AI

Unethical emotional manipulation

Realistic Risks of AI

Unethical usage: autonomous weapons?

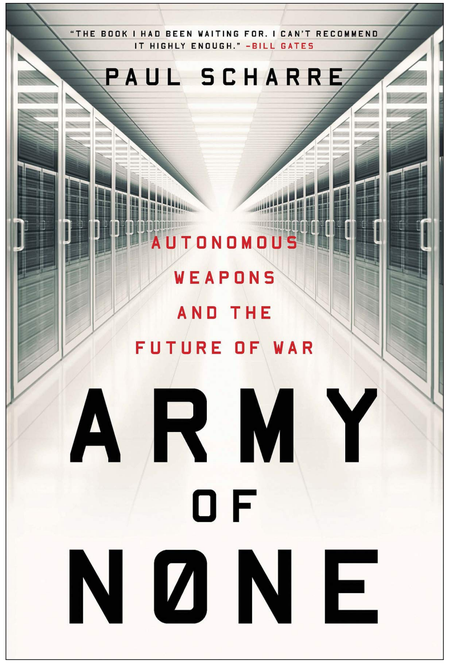

https://autonomousweapons.org/

Realistic Risks of AI

AI in the "wrong hands"

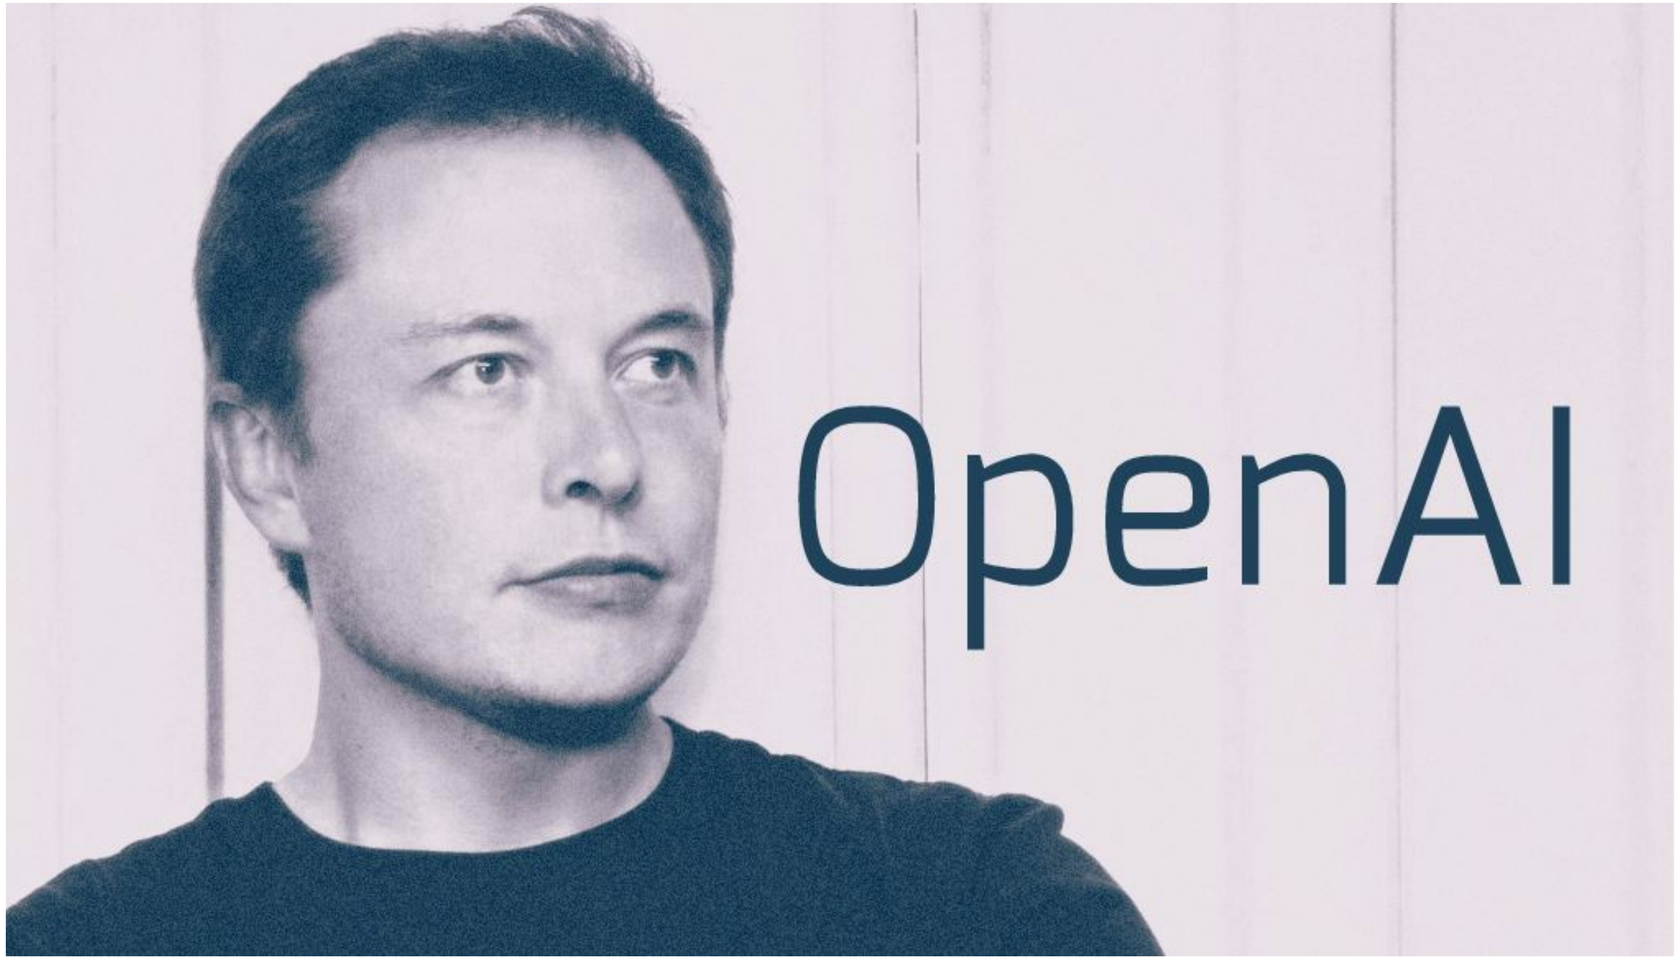

Realistic Benefits of AI

The central question:

Can we ensure that the benefits of AI outweigh the potential risks?

Realistic Benefits of AI

Significant reduction of driving fatalities

Realistic Benefits of AI

Happier, healthier lives

Realistic Benefits of AI

Increased productivity and prosperity

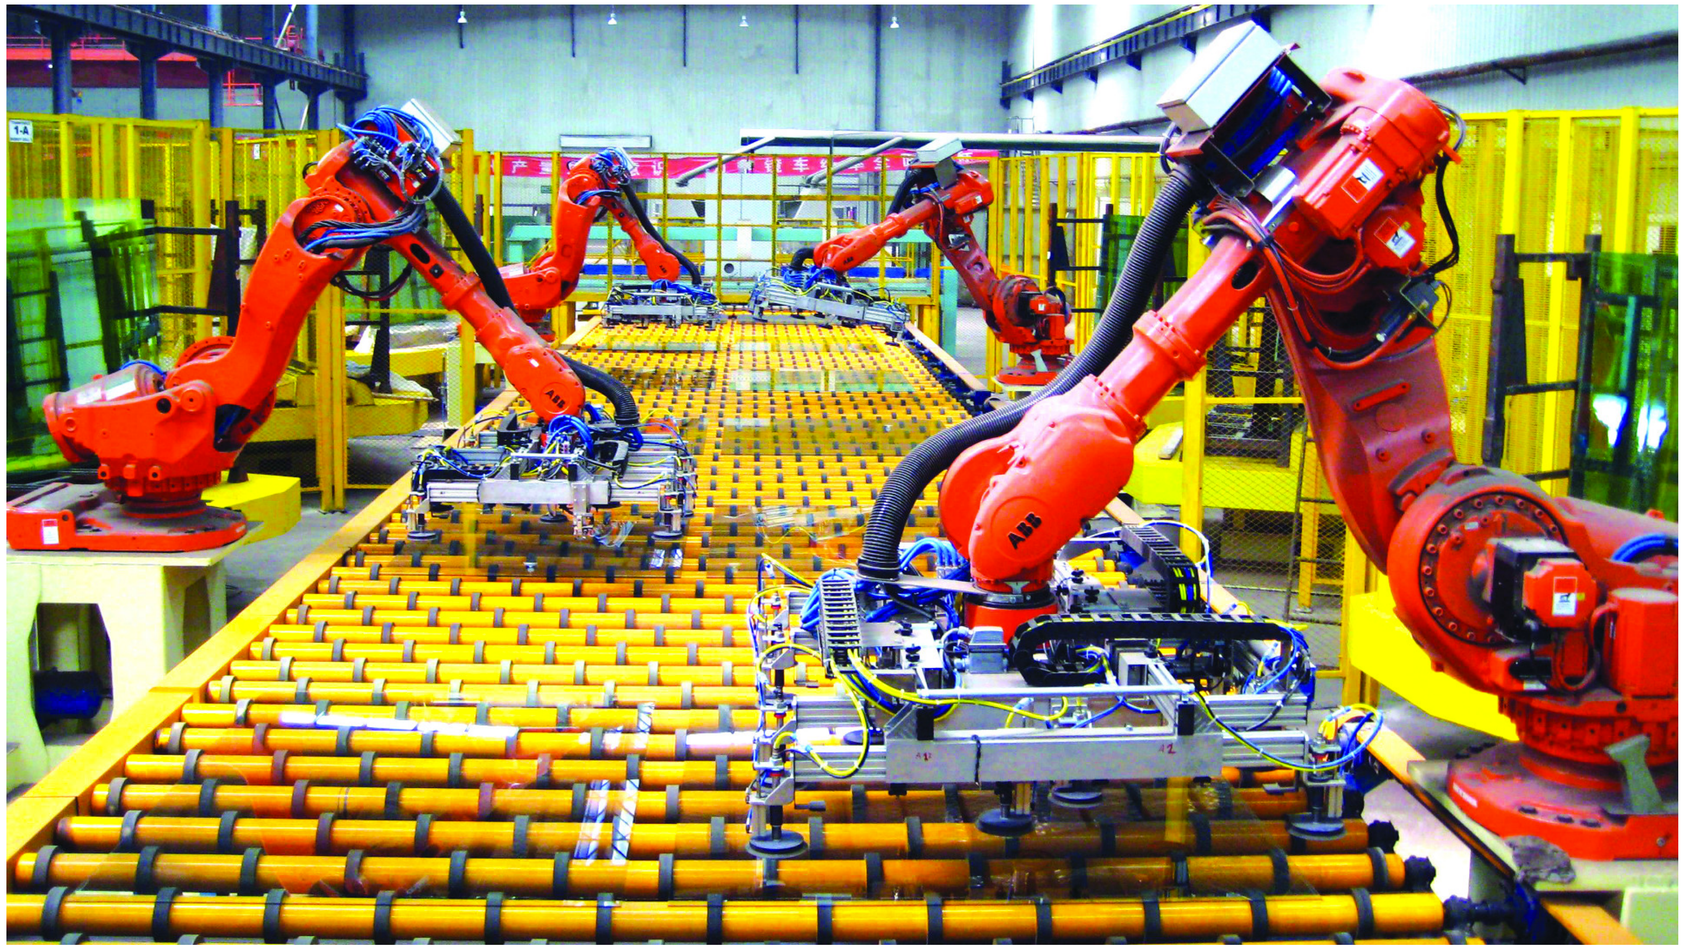

Realistic Benefits of AI

Dirty, dangerous, and dull

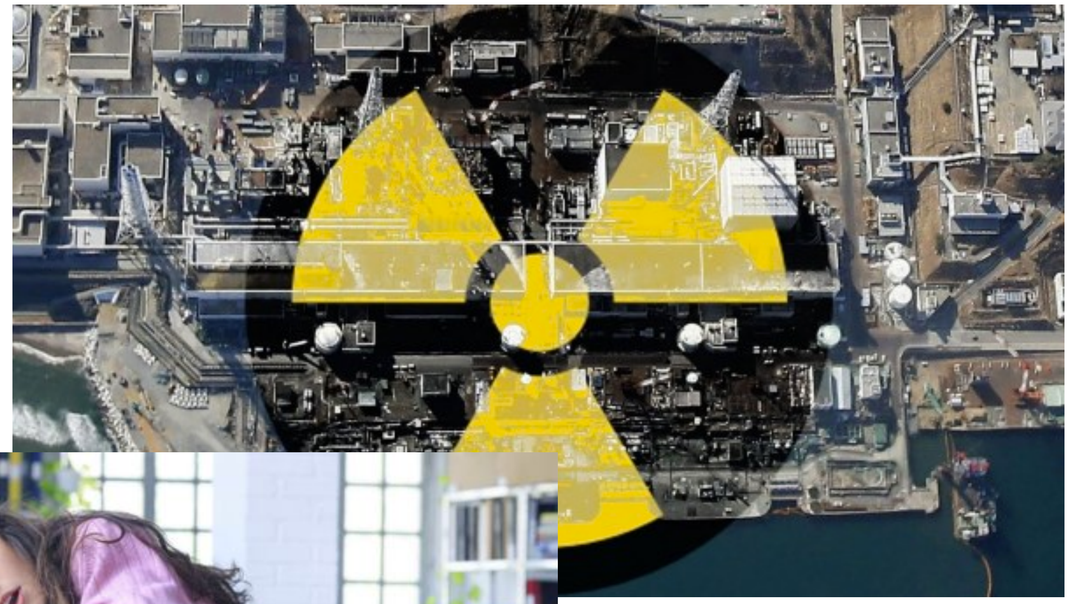

Realistic Benefits of AI

Greater social justice

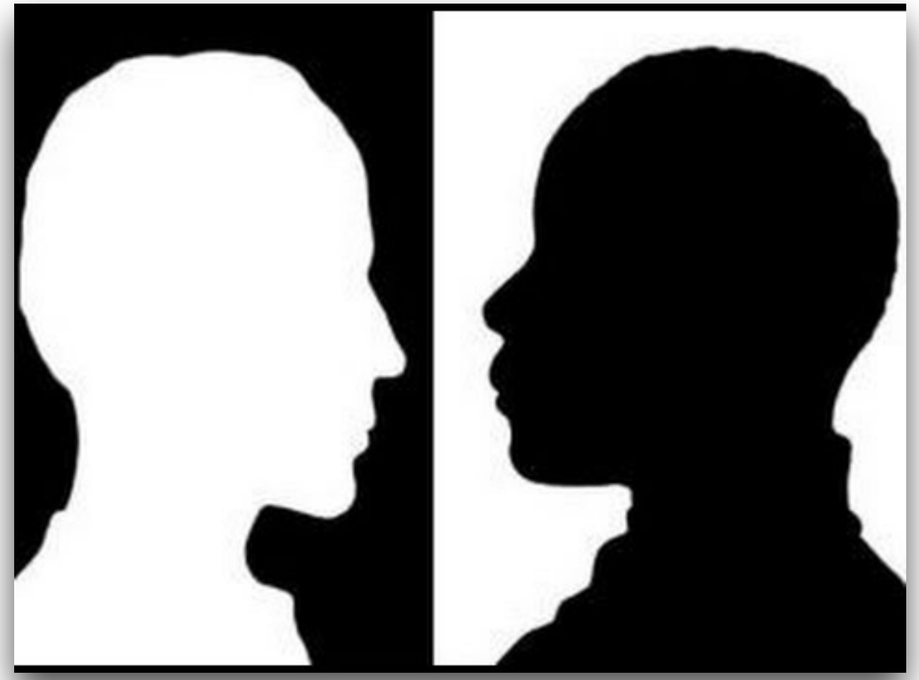

Realistic Benefits of AI

Beyond human capabilities

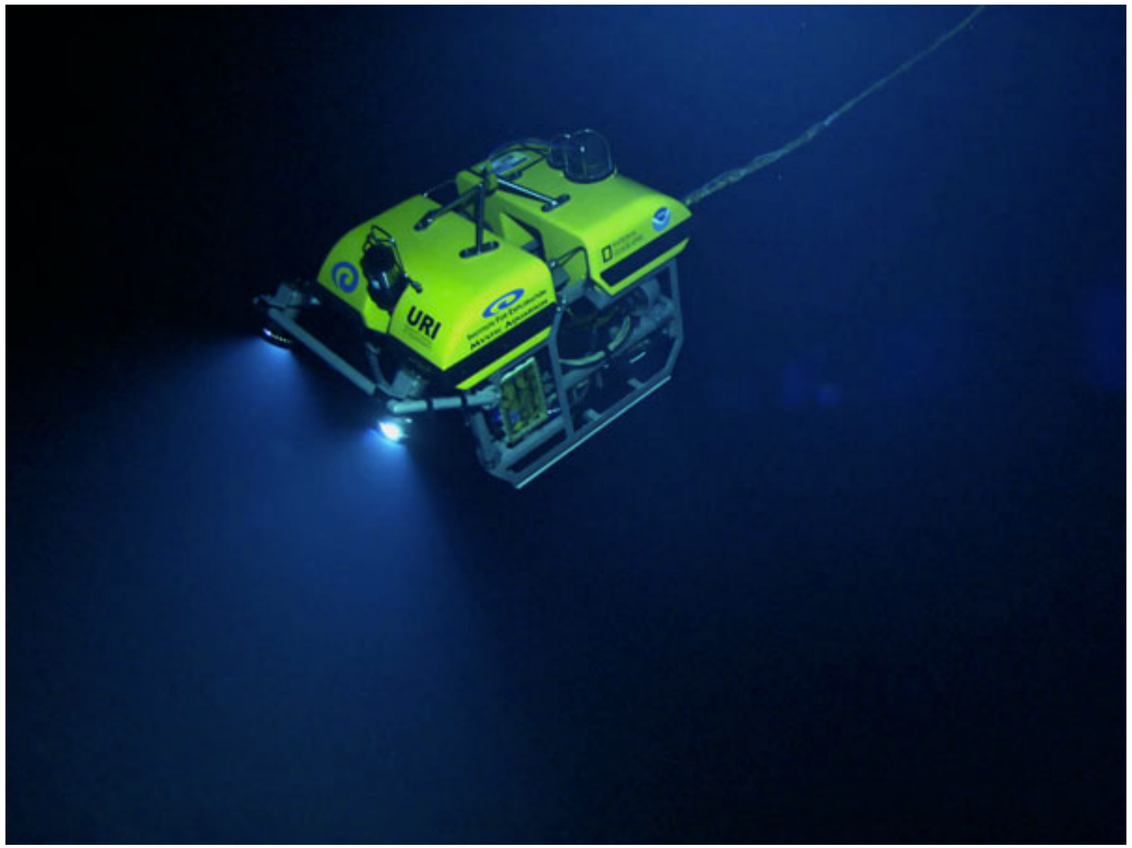

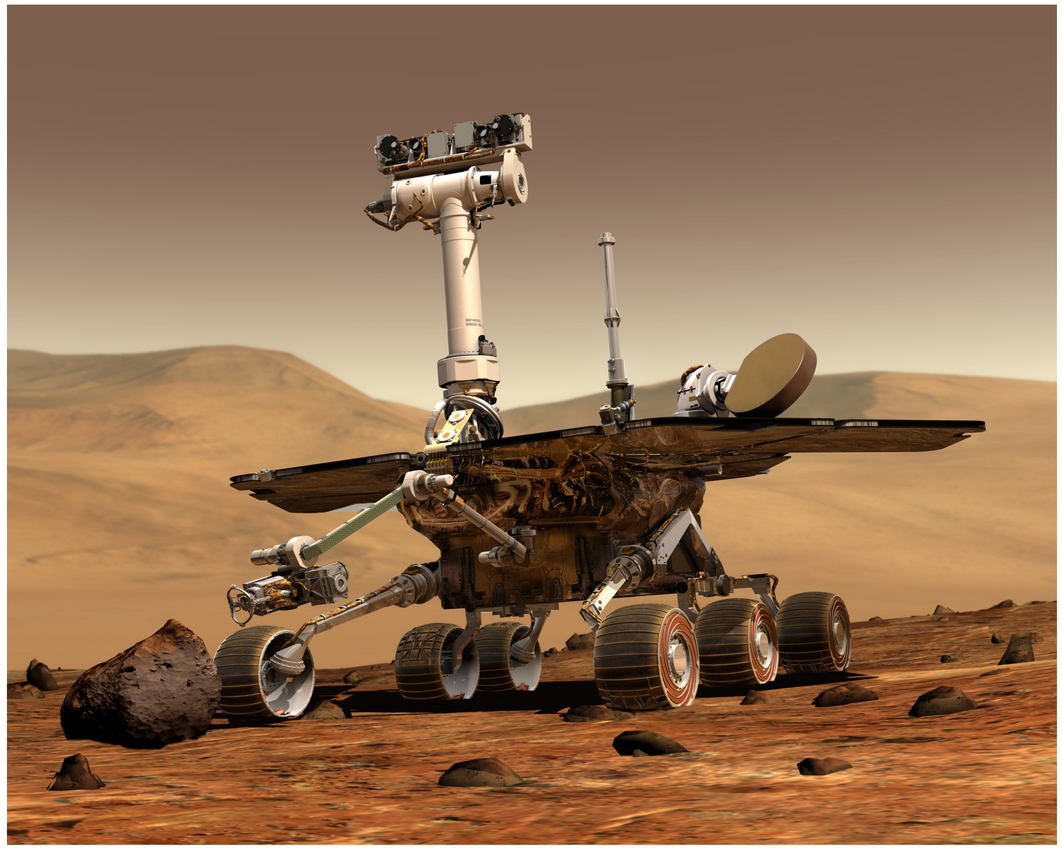

Realistic Benefits of AI

But what does social good really mean?

A Human-Centered Approach to Artificial Intelligence

Realistic Benefits of AI

But what does social good really mean?

"Artificial intelligence should **treat all people fairly**, **empower** everyone, perform **reliably** and **safely**, be **understandable**, be **secure and respect privacy**, and have **algorithmic accountability**. It should be aligned with existing **human values**, be **explainable**, be **fair**, and **respect user data rights**. It should be used for **socially beneficial purposes**, and always remain under meaningful **human control**."

— Tom Chatfield (2020)

[Source: There's No Such Thing As 'Ethical A.I.']