

Warm up

Add to the beginning of your name to one of 5 discussion categories that you would like to participate in today:

Bias

Weapons

Privacy

Liability

Jobs

Go to the following links to start the surveys for a little while...

FCE: <https://cmu.smartevals.com/>

TA survey: <https://www.ugrad.cs.cmu.edu/ta/S20/feedback/>

Announcements

Assignments

- P5
 - Due tomorrow, 4/30, 10 pm

Final Exam

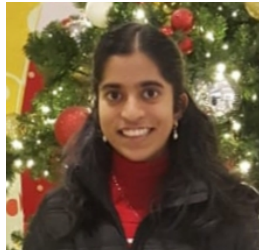
- Monday 5/4, 1-4 pm, Zoom and Gradescope
- See Piazza for details
- Recitation → Review session
- Practice exam and learning objectives out!
 - Practice exam solution session

FCE: <https://cmu.smartevals.com/>

TA survey: <https://www.ugrad.cs.cmu.edu/ta/S20/feedback/>



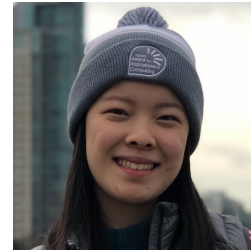
Claire (Head TA)



Aarthi



Ahana



Amy



Harlene



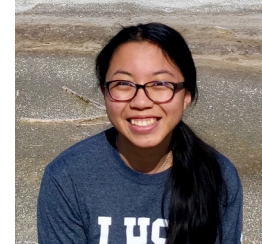
Maia



Sean



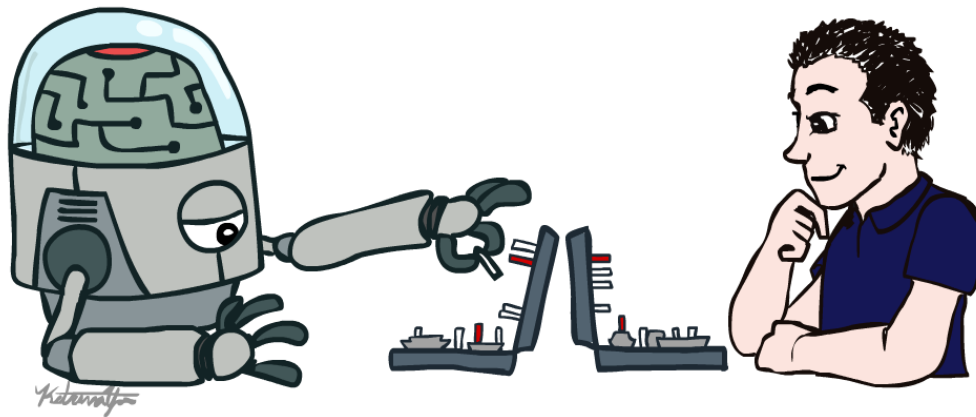
Simran



Tina

AI: Representation and Problem Solving

AI and Ethical Decision Making



Instructors: Pat Virtue & Stephanie Rosenthal

Slide credits: CMU AI and <http://ai.berkeley.edu>

Intelligent Agents

From first lecture

Candy Grab

- A. 11 pieces on the table
- B. Take turns taking either 1 or 2 pieces
- C. Person that takes the last piece wins!

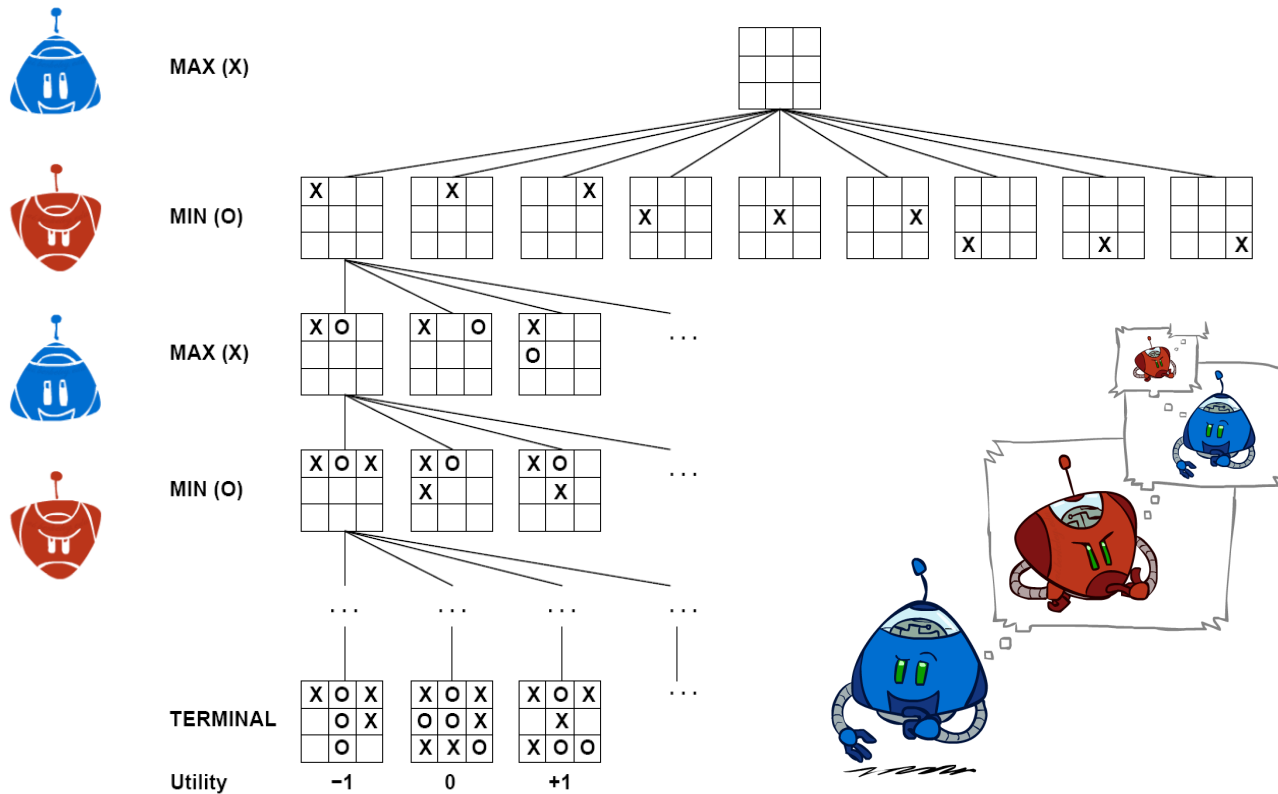
```
class Agent
    function getAction(state)
        return action
```

Piazza Poll 1

Three “Intelligent” Agents

Which agent code is the most “intelligent”?

A: Search



Games – Three “Intelligent” Agents

B: Encode the pattern

```
function getAction( numPiecesAvailable )  
  
    if numPiecesAvailable % 3 == 2  
        return 2  
    else  
        return 1
```

```
10's value:Win  
9's value:Lose  
8's value:Win  
7's value:Win  
6's value:Lose  
5's value:Win  
4's value:Win  
3's value:Lose  
2's value:Win  
1's value:Win  
0's value:Lose
```

Games – Three “Intelligent” Agents

C: Record statistics of winning positions

Pieces Available	Take 1	Take 2
2	0%	100%
3	2%	0%
4	75%	2%
5	4%	68%
6	5%	6%
7	60%	5%

Piazza Poll 1

Three “Intelligent” Agents

Which agent code is the most “intelligent”?

- A. Search
- B. Encode multiple of 3 pattern
- C. Keep stats on winning positions

Piazza Poll 2

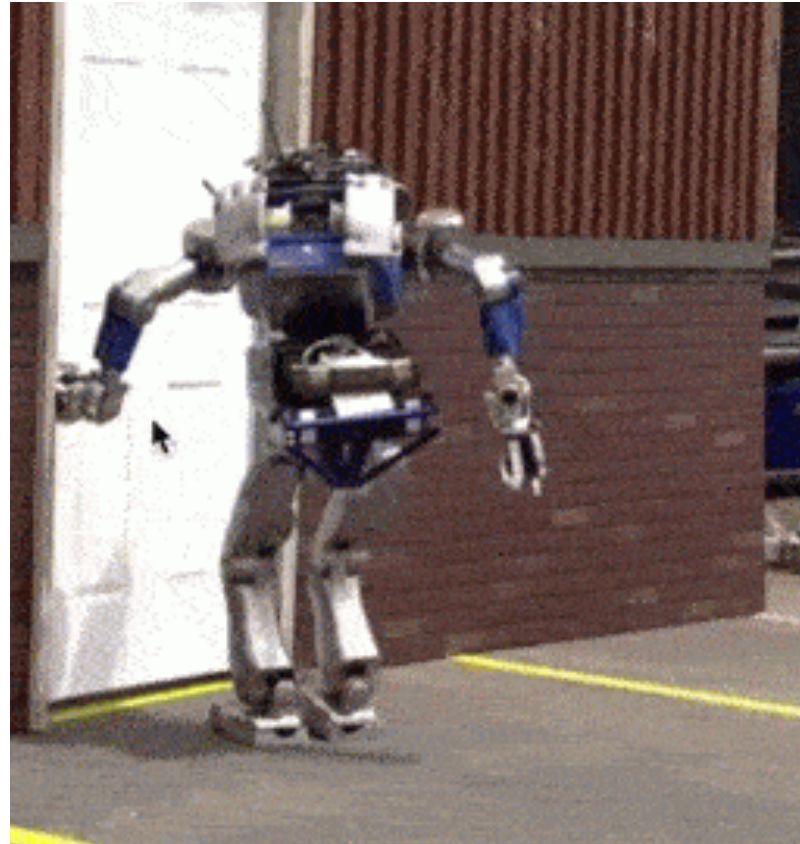
Which 281 technique is the most intelligent?

- A. Search
- B. Logical inference
- C. Numeric optimization
- D. Q-learning
- E. Approximate Q-learning
- F. Exact inference Bayes nets
- G. Approximate inference Bayes nets

State of the Art in AI

Should We Worry about Today's A.I.?

https://www.youtube.com/watch?v=UUOo8N9_iH0



AI in the News

AI in the News



Nuro Driverless Vehicles Become Robot Pack Mules For California COVID-19 Medical Centers

By Alan Ohnsman Forbes Staff

The Silicon Valley startup is deploying some of its small robotic delivery vehicles to haul critical supplies for medical staff caring for coronavirus patients.

<https://www.forbes.com/sites/alanohnsman/2020/04/22/nuro-driverless-vehicles-become-robot-pack-mules-for-california-covid-19-medical-centers>

EDITORS' PICK | 876 views | Apr 22, 2020, 06:32pm EDT

Einride Set To Bring Electric Driverless Trucks To Grocery Stores In Sweden, Then The World



Sebastian Blanco Contributor 
Transportation

f

t

in



Einride's Autonomous Pod Truck with Carl Ceder, Logistics Manager of Lidl Sweden. EINRIDE

<https://www.forbes.com/sites/sebastianblanco/2020/04/22/einride-ready-to-bring-electric-trucks-to-lidls-swedish-operations-then-the-world>

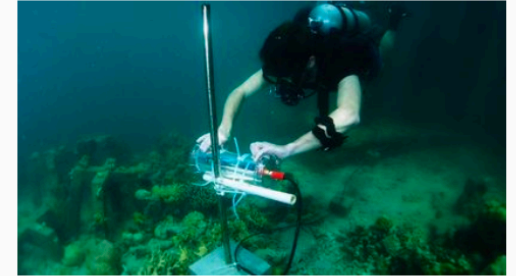
AI in the News

Apr 22, 2020

Intel & Accenture Use AI To Save The Coral Reef

Accenture, Intel and the Sulubaaï Environmental Foundation announced Project: CORaiL, an artificial intelligence (AI) - powered solution to monitor, characterize and analyze coral reef resilience.

By **Jun Wu** Contributor 



Apr 22, 2020

How Artificial Intelligence, IoT And Big Data Can Save The Bees

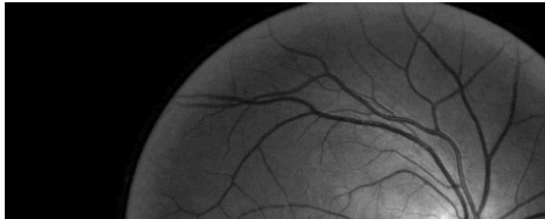
Artificial intelligence (AI), internet of things sensors, big data analytics, and other technologies are being used to monitor and develop solutions to mitigate the global issue of declining honey bee and pollinator populations.

By **Bernard Marr** Contributor 



AI in the News

MIT
Technology
Review



Artificial Intelligence / Machine learning

Google's medical AI was super accurate in a lab. Real life was a different story.

If AI is really going to make a difference to patients we need to know how it works when real humans get their hands on it, in real situations.

by Will Douglas Heaven

April 27, 2020

439 views | Apr 23, 2020, 11:30am EDT

Dangers Await Self-Driving Car Tryouts As Post-Pandemic Roadway Efforts Get Underway



Lance Eliot Contributor

Transportation

Dr. Lance B. Eliot is a world-renowned expert on Artificial Intelligence (AI) and Machine Learning (ML).



237 views | Apr 21, 2020, 12:15pm EDT

Low-Trust Versus High-Trust Beliefs Will Impact The Growth Of AI, Including For Self-Driving Cars



Lance Eliot Contributor

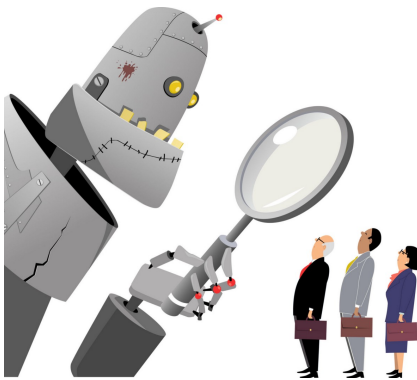
Transportation

Dr. Lance B. Eliot is a world-renowned expert on Artificial Intelligence (AI) and Machine Learning (ML).



Should We Worry about Today's A.I.?

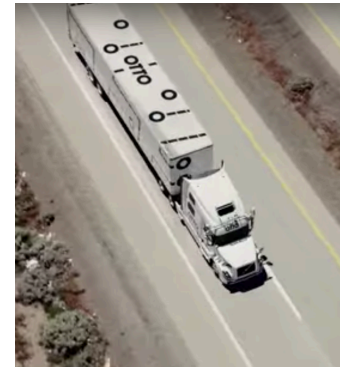
Bias



Privacy



Jobs



Weapons



Liability



Bias

How are AI systems motivated? How does this motivation cause bias?

How could we measure bias (including cultural bias) in an AI system?

How do we as programmers introduce bias into the AI techniques we've learned about in 15-281?

What are some ways that bias in AI programming can impact the world?

Domains to think about:

- Self-driving cars

- Hiring and judicial decision making

- Traffic control

Weapons

Should AI be used as a weapon or in warfare?

Where would you draw the line for how it could be used and how not?

Can you think of some rules or laws that may help prevent unsafe AI?

Domains to think about:

- Social media bots and information warfare

- Drones and autonomous killing machines

- Data collection and surveillance technologies

Privacy

What are some tradeoffs we make when giving AI our private information?
Do you have a limit for the type or amount of data you are willing to provide?
How can we make the data usage of AI systems more explicit?
Are there ways that we can still enjoy the benefits of AI systems without giving up our private information?

Domains to think about:

Social media

Search (Google, Amazon, Yelp, ...)

Healthcare

Liability

How do we determine what went wrong (or right) in deployed AI systems?

Which techniques that we learned about have explainable results?

Who is responsible for “mistakes” that AI systems make?

What if someone or another AI misuses AI technology? Who is responsible?

Once autonomous vehicles are readily available, should it be illegal for humans to drive?

Domains to think about:

- Self-driving cars

- Healthcare decision-making

- Land and Air traffic control

Jobs

What types of jobs should AI do instead of people?

Is it ok if autonomous vehicles completely replace human drivers?

What will people do if their jobs are taken?

How can AI systems improve or change people's jobs even if it doesn't take the jobs?

Domains to think about:

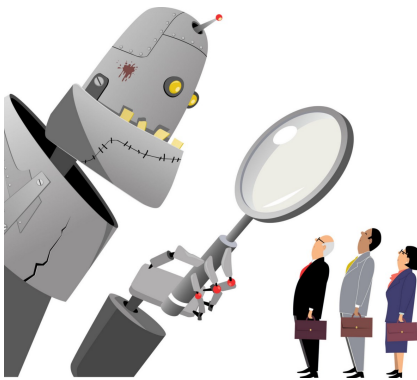
- Driving and Logistics

- Healthcare decision-making

- Entertainment

Should We Worry about Today's A.I.?

Bias



Privacy



Jobs



Weapons



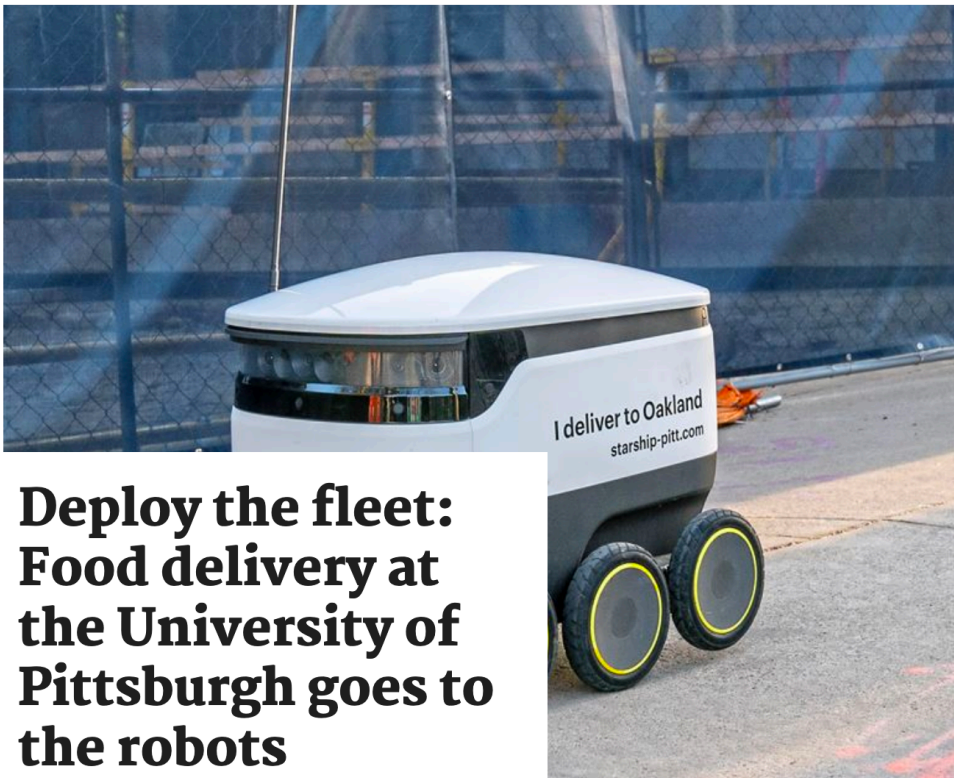
Liability



Additional Resources and Ideas

AI Bias

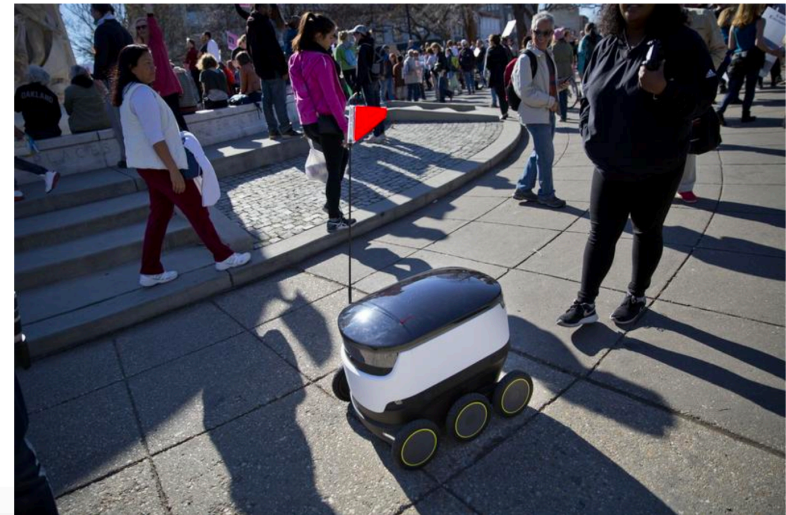
Question: What technique could have been used to learn from the mistake?



Food Delivery Robots Pulled From Pitt Campus After Backlash About Mobility

By KATHLEEN J. DAVIS • OCT 22, 2019

[Share](#) [Tweet](#) [Email](#)



AI Bias

5 Examples Of How AI Can Be Used Across The Supply Chain Forbes



Blake Morgan Senior Contributor

CMO Network

I am a Customer Experience Futurist, Author and Keynote Speaker.



Search/Navigation:

Create efficient routes for the fleet of trucks

Use of robots to deliver medicine, groceries etc

CSPs:

Deciding which stores get preference for deliveries

Deciding what products go to what stores

AI Bias

AI could be the key to ending discrimination in hiring, but experts warn it can be just as biased as humans

Aaron Holmes Oct 8, 2019, 12:14 PM



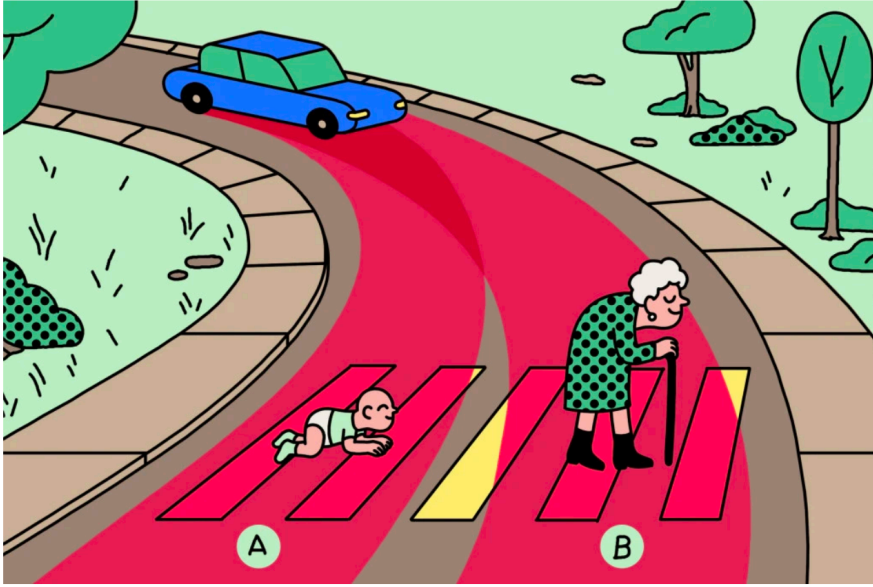
Federal study confirms racial bias of many facial-recognition systems, casts doubt on their expanding use

Asian and African American people were up to 100 times more likely to be misidentified than white men, depending on the particular algorithm and type of search. Native Americans had the highest false-positive rate of all ethnicities, according to the study, which found that systems varied widely in their accuracy.

The faces of African American women were falsely identified more often in the kinds of searches used by police investigators where an image is compared to thousands or millions of others in hopes of identifying a suspect.

Algorithms developed in the United States also showed high error rates for “one-to-one” searches of Asians, African Americans, Native Americans and Pacific Islanders. Such searches are critical to functions including cellphone sign-ons and airport boarding schemes, and errors could make it easier for impostors to gain access to those systems.

<https://www.washingtonpost.com/technology/2019/12/19/federal-study-confirms-racial-bias-many-facial-recognition-systems-casts-doubt-their-expanding-use/>



SIMON LANDREIN

Tech policy / AI Ethics

Should a self-driving car kill the baby or the grandma? Depends on where you're from.

The infamous "trolley problem" was put to millions of people in a global study, revealing how much ethics diverge across cultures.

TECHNOLOGY

The Ethics of Autonomous Cars

Sometimes good judgment can compel us to act illegally. Should a self-driving vehicle get to make that same decision?

PATRICK LIN OCTOBER 8, 2013

NEWS · 24 OCTOBER 2018

Self-driving car dilemmas reveal that moral choices are not universal

Survey maps global variations in ethics for programming autonomous vehicles.

<https://www.nature.com/articles/d41586-018-07135-0#ref-CR1>

<https://www.technologyreview.com/2018/10/24/139313/a-global-ethics-study-aims-to-help-ai-solve-the-self-driving-trolley-problem/>

<https://www.theatlantic.com/technology/archive/2013/10/the-ethics-of-autonomous-cars/280360/>

AI Bias



Alexandra Chouldechova
CMU, Statistics and Public Policy
<http://www.contrib.andrew.cmu.edu/~achoulde/>

<https://fatconference.org/>

AI Weapons



<https://fatconference.org/>

AI Weapons



8,353 views | Aug 1, 2019, 02:49am EDT

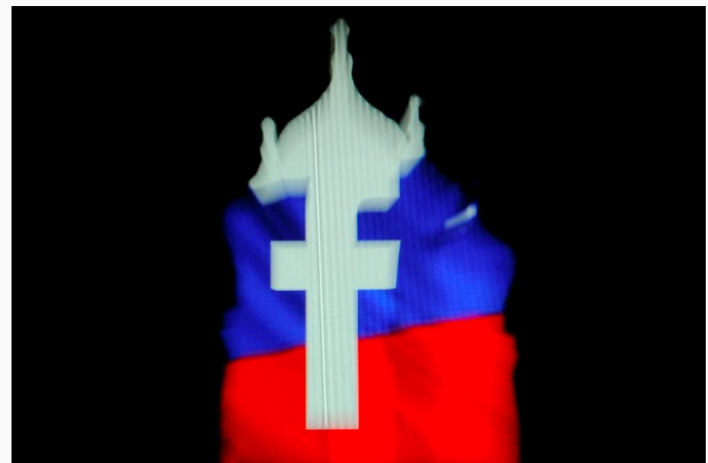
Cyber Warfare: Army Deploys 'Social Media Warfare' Division To Fight Russia



Zak Doffman Contributor

[Cybersecurity](#)

I write about security and surveillance.

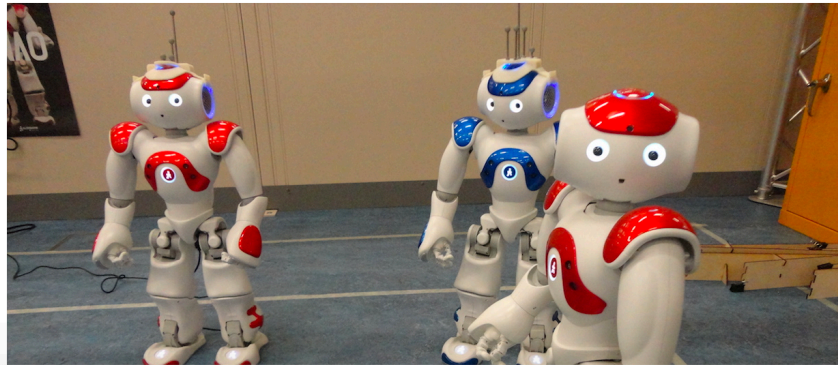


06 Feb 2018 | 21:28 GMT

Why Ethical Robots Might Not Be Such a Good Idea After All

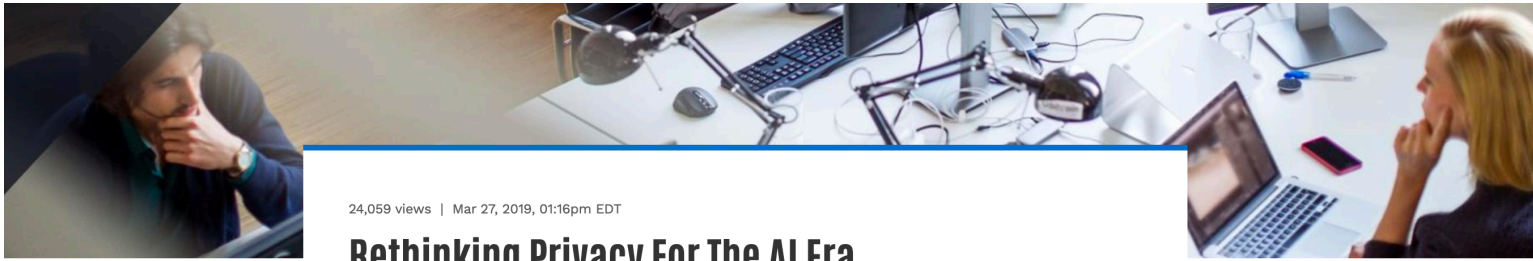
The risks that a robot's ethics might be compromised by unscrupulous actors raise doubts over the wisdom of embedding ethical decision making in real-world safety critical robots

By **Alan Winfield**



<https://spectrum.ieee.org/automaton/robotics/artificial-intelligence/why-ethical-robots-might-not-be-such-a-good-idea-after-all>

AI Privacy



24,059 views | Mar 27, 2019, 01:16pm EDT

Rethinking Privacy For The AI Era

Insights Team Insights Contributor
Forbes **INSIGHTS** With **Intel AI** | Paid Program
[Innovation](#)

[f](#) [t](#) [in](#)

Concerns over consumer privacy have peaked in recent years—roughly in step with the rise of advanced technologies like artificial intelligence. About 9 in 10 American internet users say they are concerned about the privacy and security of their personal information online, and 67% are now advocating for strict national privacy laws, according to a [study by](#) Intouch International.

Fed up by a steady stream of incidents that range from the 2017 [Equifax hack](#) to the nefarious gaming of consumers' social media data for political purposes, policymakers have begun to strike back on consumers' behalf.

<https://www.forbes.com/sites/insights-intelai/2019/03/27/rethinking-privacy-for-the-ai-era>

AI Liability

<https://undark.org/2019/12/04/black-box-artificial-intelligence/>

NEWS

Unpacking the Black Box in Artificial Intelligence for Medicine

Deep learning will radically change aspects of our medical care. How well do we need to understand how AI tools work?

Visual: Yuichiro Chino / Getty Images



AI Jobs

<https://www.supplychaindigital.com/technology/ups-invests-autonomous-driving-firm-tusimple>

SUPPLY CHAIN

LOGISTICS

TECHNOLOGY

UPS invests in autonomous driving firm TuSimple

By SEAN GALEA-PACE • Aug 15, 2019, 10:20AM



AI Jobs

Is it ok if autonomous vehicles completely replace human drivers?

- <https://www.brookings.edu/research/what-jobs-are-affected-by-ai-better-paid-better-educated-workers-face-the-most-exposure/>
- <https://www.vox.com/platform/amp/policy-and-politics/2019/12/3/20965464/2020-presidential-candidates-jobs-automation-ai>
- <https://www.irishtimes.com/business/technology/short-window-to-stop-ai-taking-control-of-society-warns-ex-google-employee-1.4104535>

AI Jobs

Should Robots or People Do These Jobs?

A Survey of Robotics Experts and Non-Experts About Which Jobs Robots Should Do

Wendy Ju & Leila Takayama
Willow Garage
68 Willow Road
Menlo Park, California, USA 94025
[wendyju, takayama] @ willowgarage.com

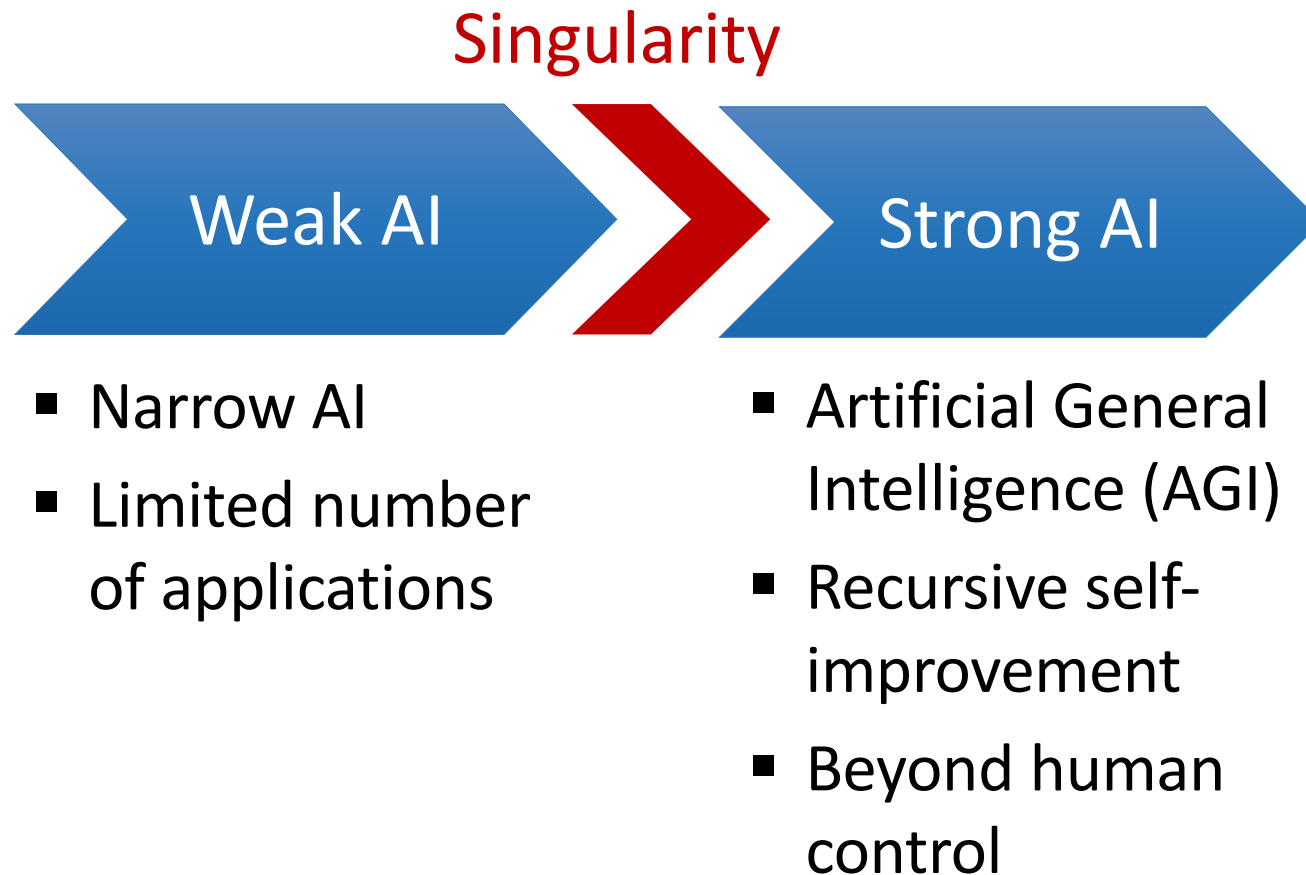
http://www.wendyju.com/publications/ju_iros11.pdf

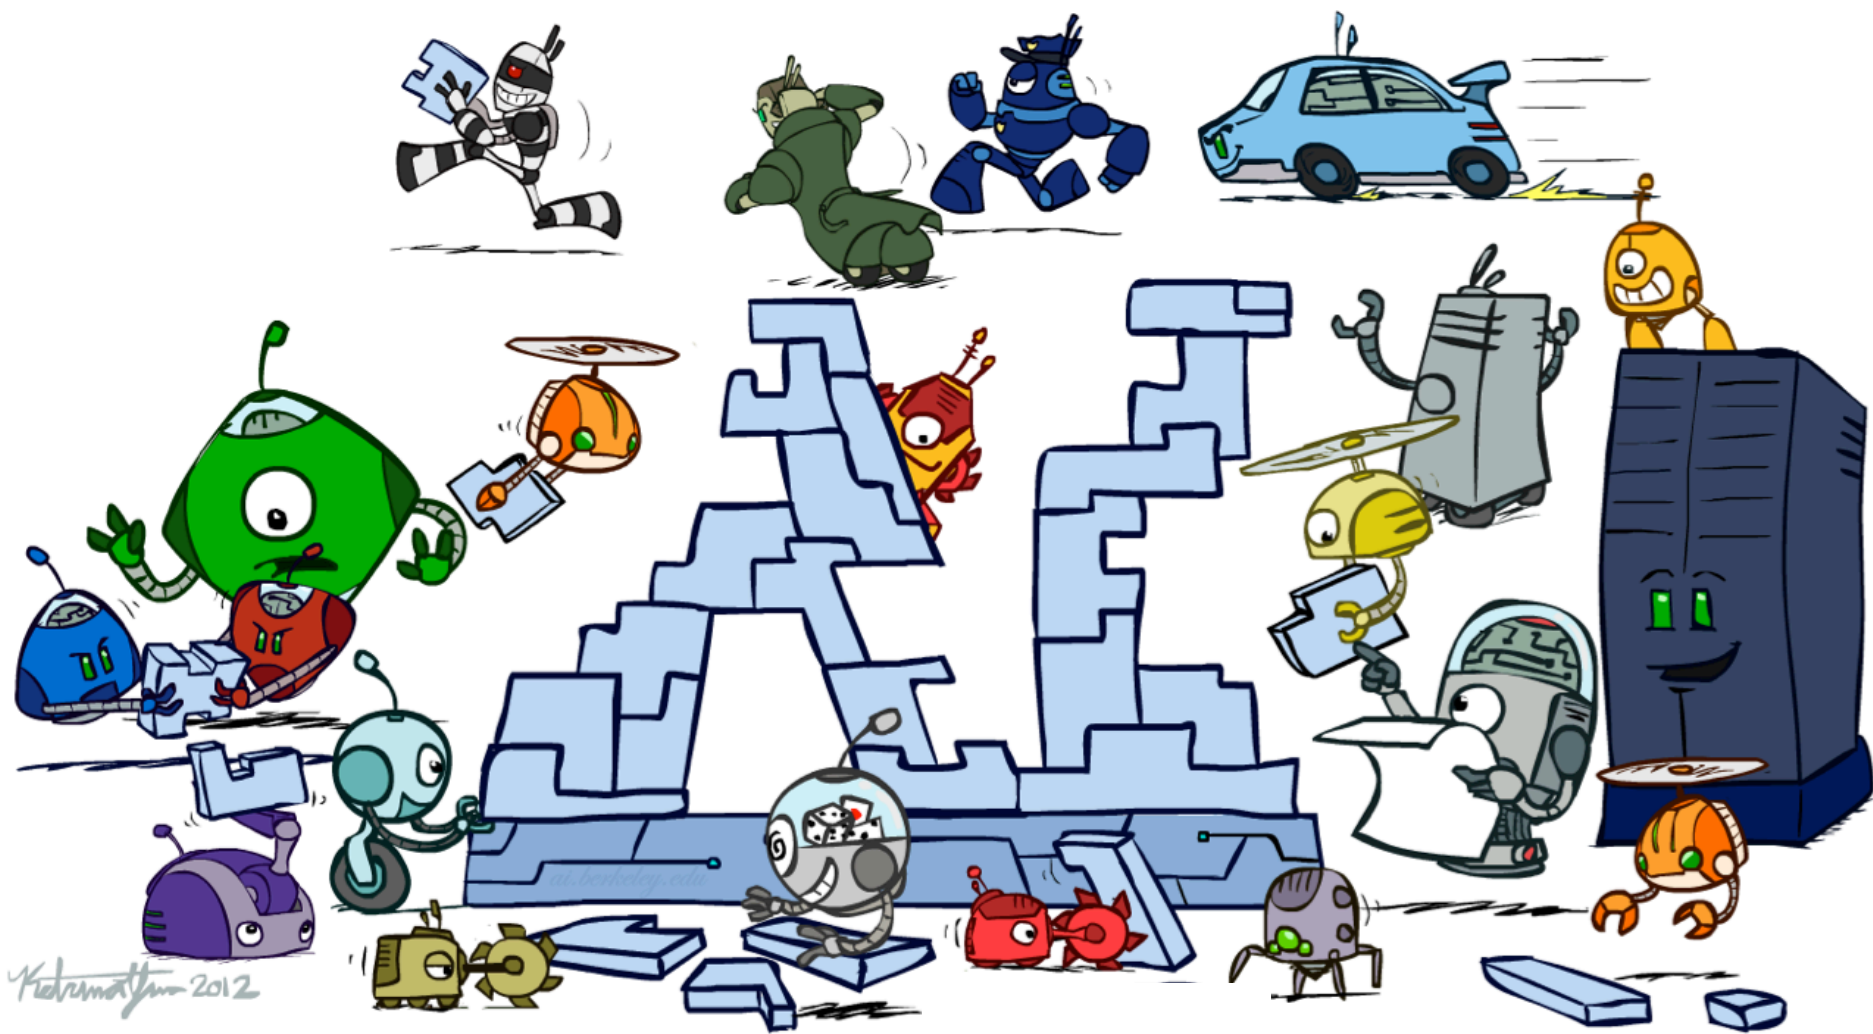
Experts SHOULD	Guiding, Directing, and Motivating Subordinates	.23***
	Speech Clarity (Importance)	.17*
	Thinking Creatively (Importance)	.17**
	Category Flexibility (Importance)	.16**
	Dependability	.14**
	Building and Construction Knowledge (Importance)	.08*
	Interpreting the Meaning of Information (Importance)	-.31***
	Cramped Workspace and Awkward Positions	-.17***
	Negotiation (Level)	-.14*
	Multi-limb Coordination (Level)	-.13*
	Initiative	-.13**
	Artistic Interests	-.11**

* $p < .05$, ** $p < .01$, *** $p < .001$. Shaded rows indicate factors where robots are preferred over people; in unshaded rows, people are preferred.

What about the Future?

Should we worry about future A.I.?





Thanks to Our Course Staff!!

Administrative Assistant

Liv Zane

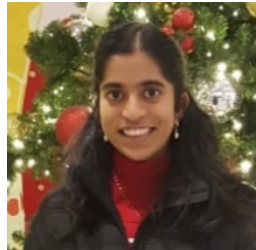
Production Crew

Pat and Stephanie!

Teaching Assistants



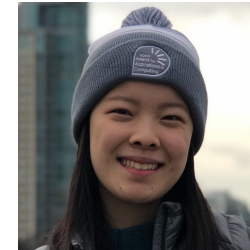
Claire (Head TA)



Aarthi



Ahana



Amy



Harlene



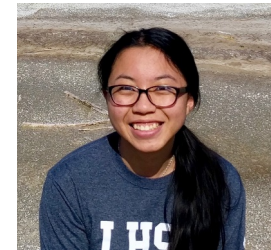
Maia



Sean



Simran



Tina