Societal Impact of ML

Machine Learning: A success story

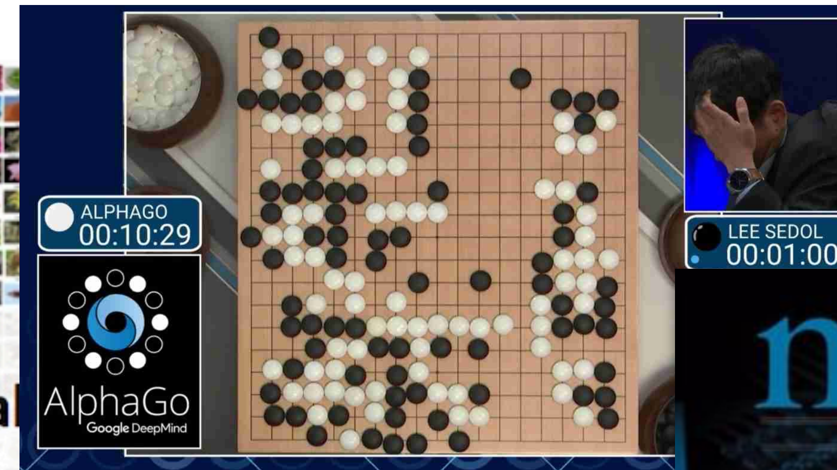

http://www.worldtonetech.com/

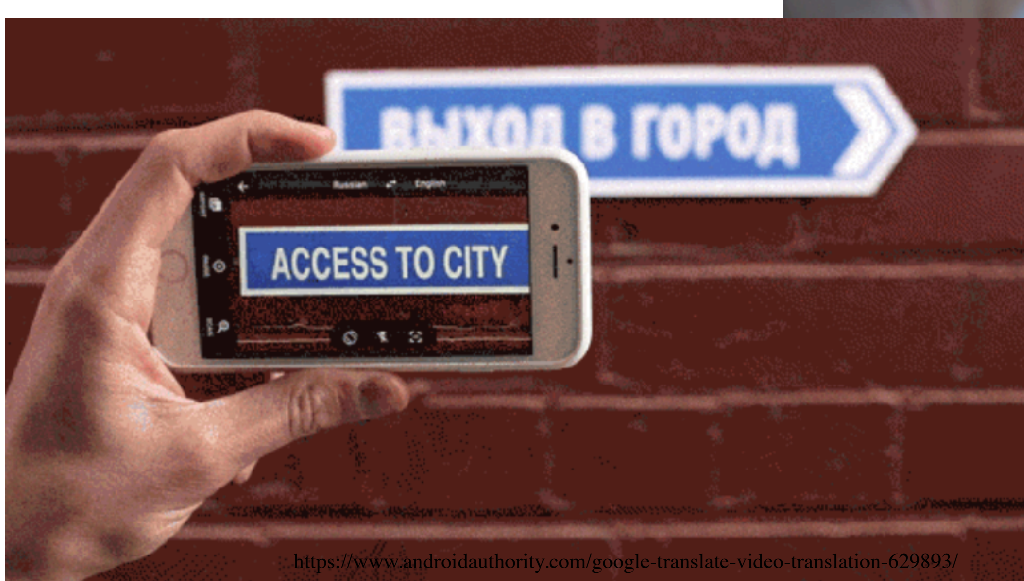

https://www.androidauthority.com/google-translate-video-translation-629893/

https://medium.com/ai-musings/the-tyranny-of-appearances-in-ai-ef3c34515711

The world is ready for ML

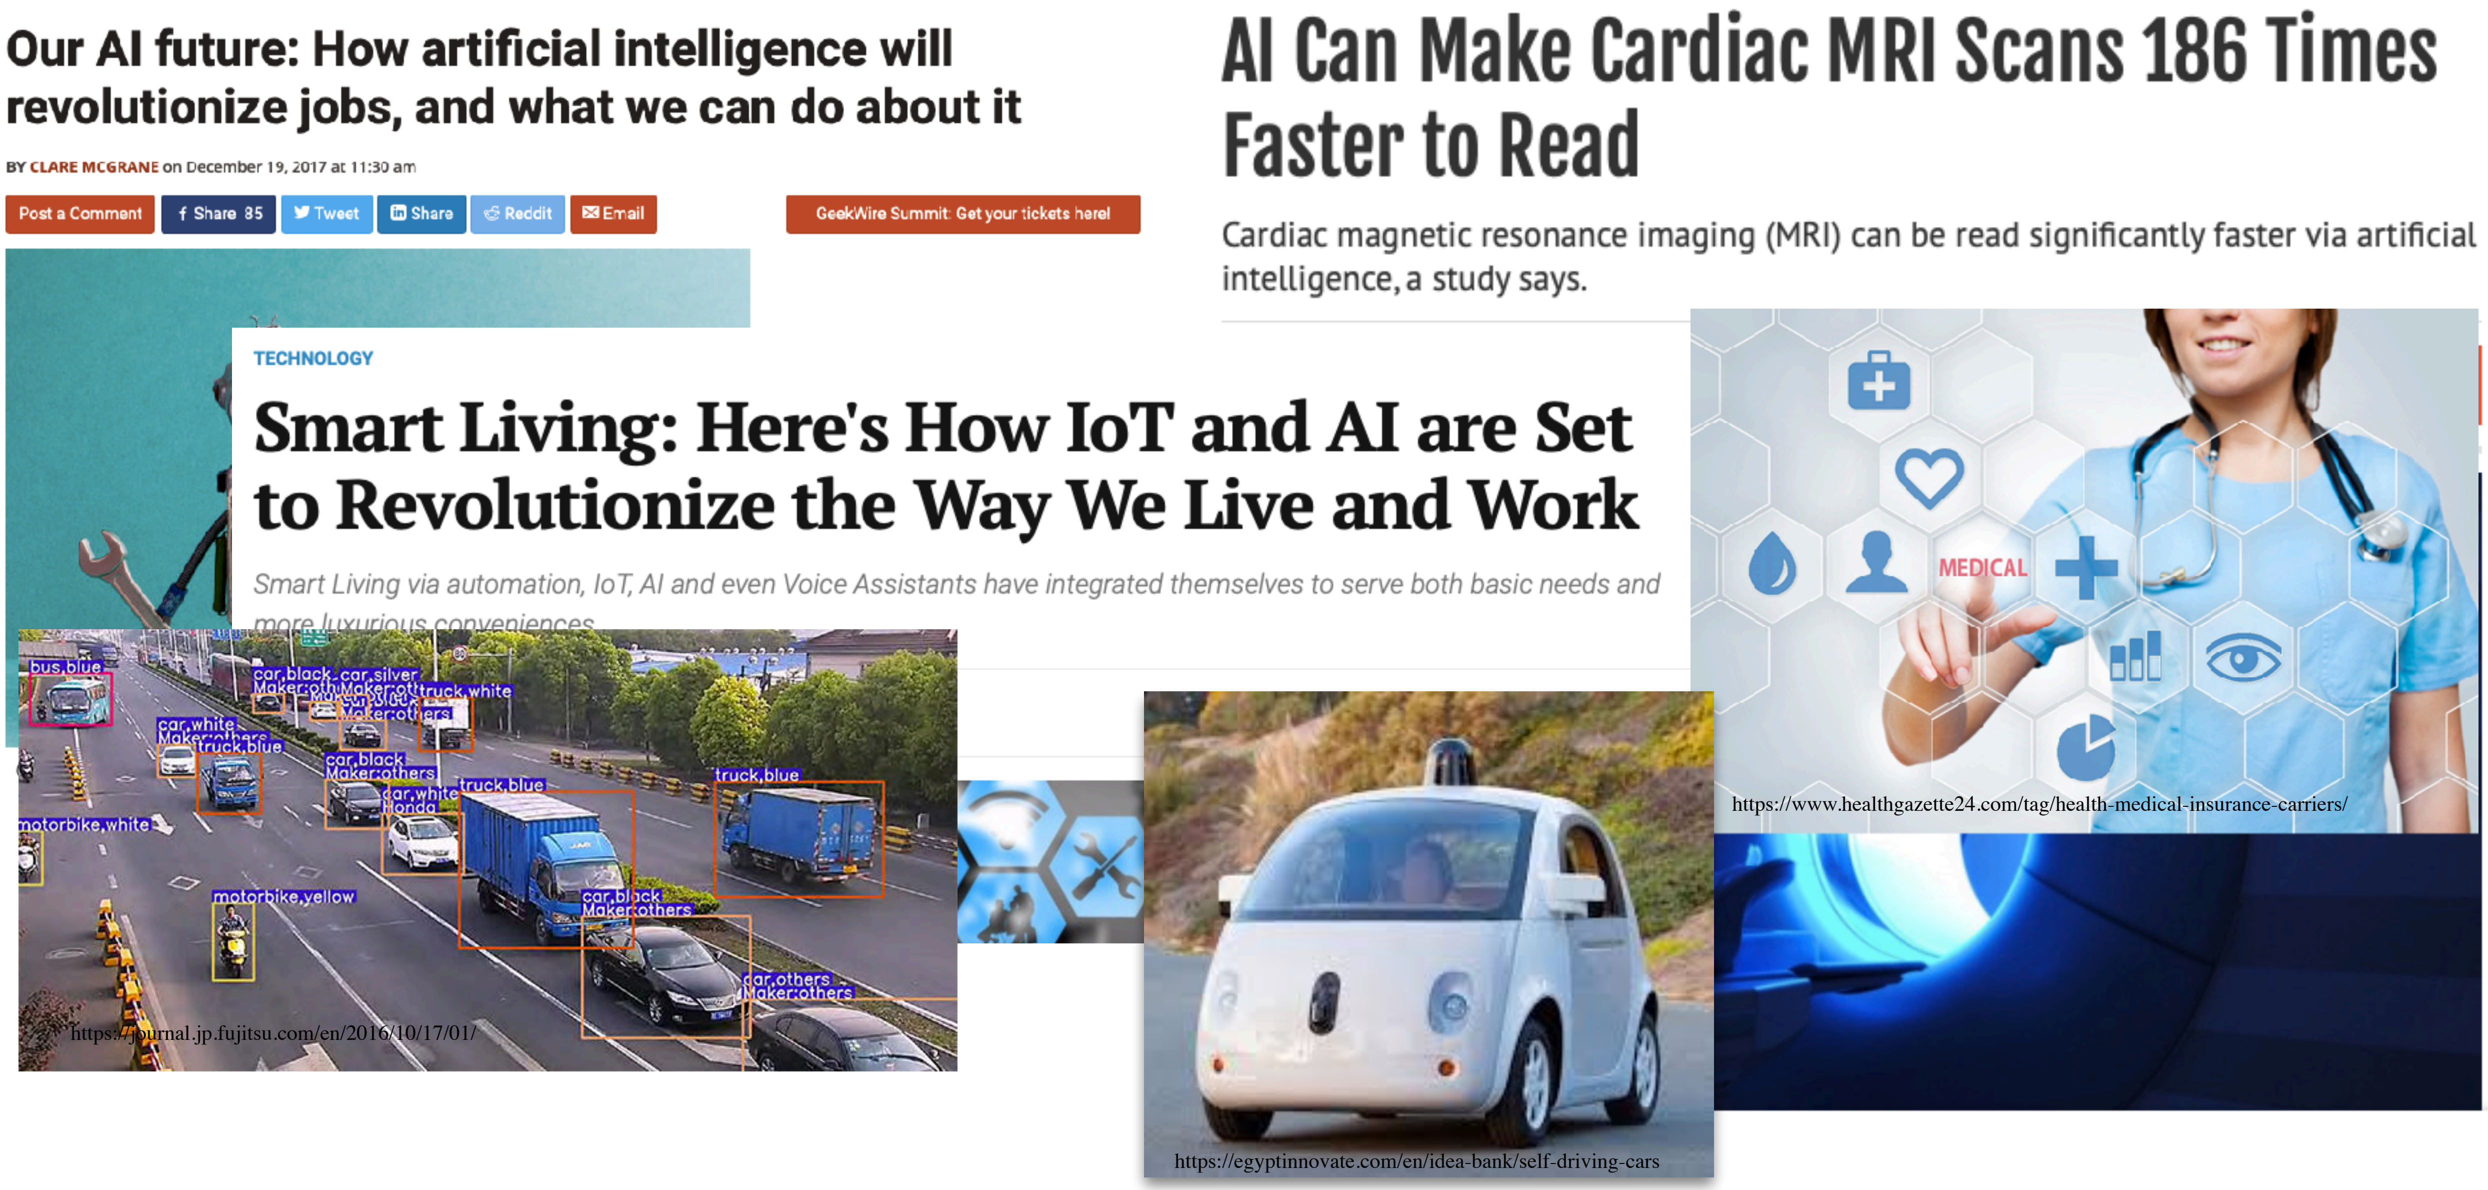

… but is ML ready for the world?

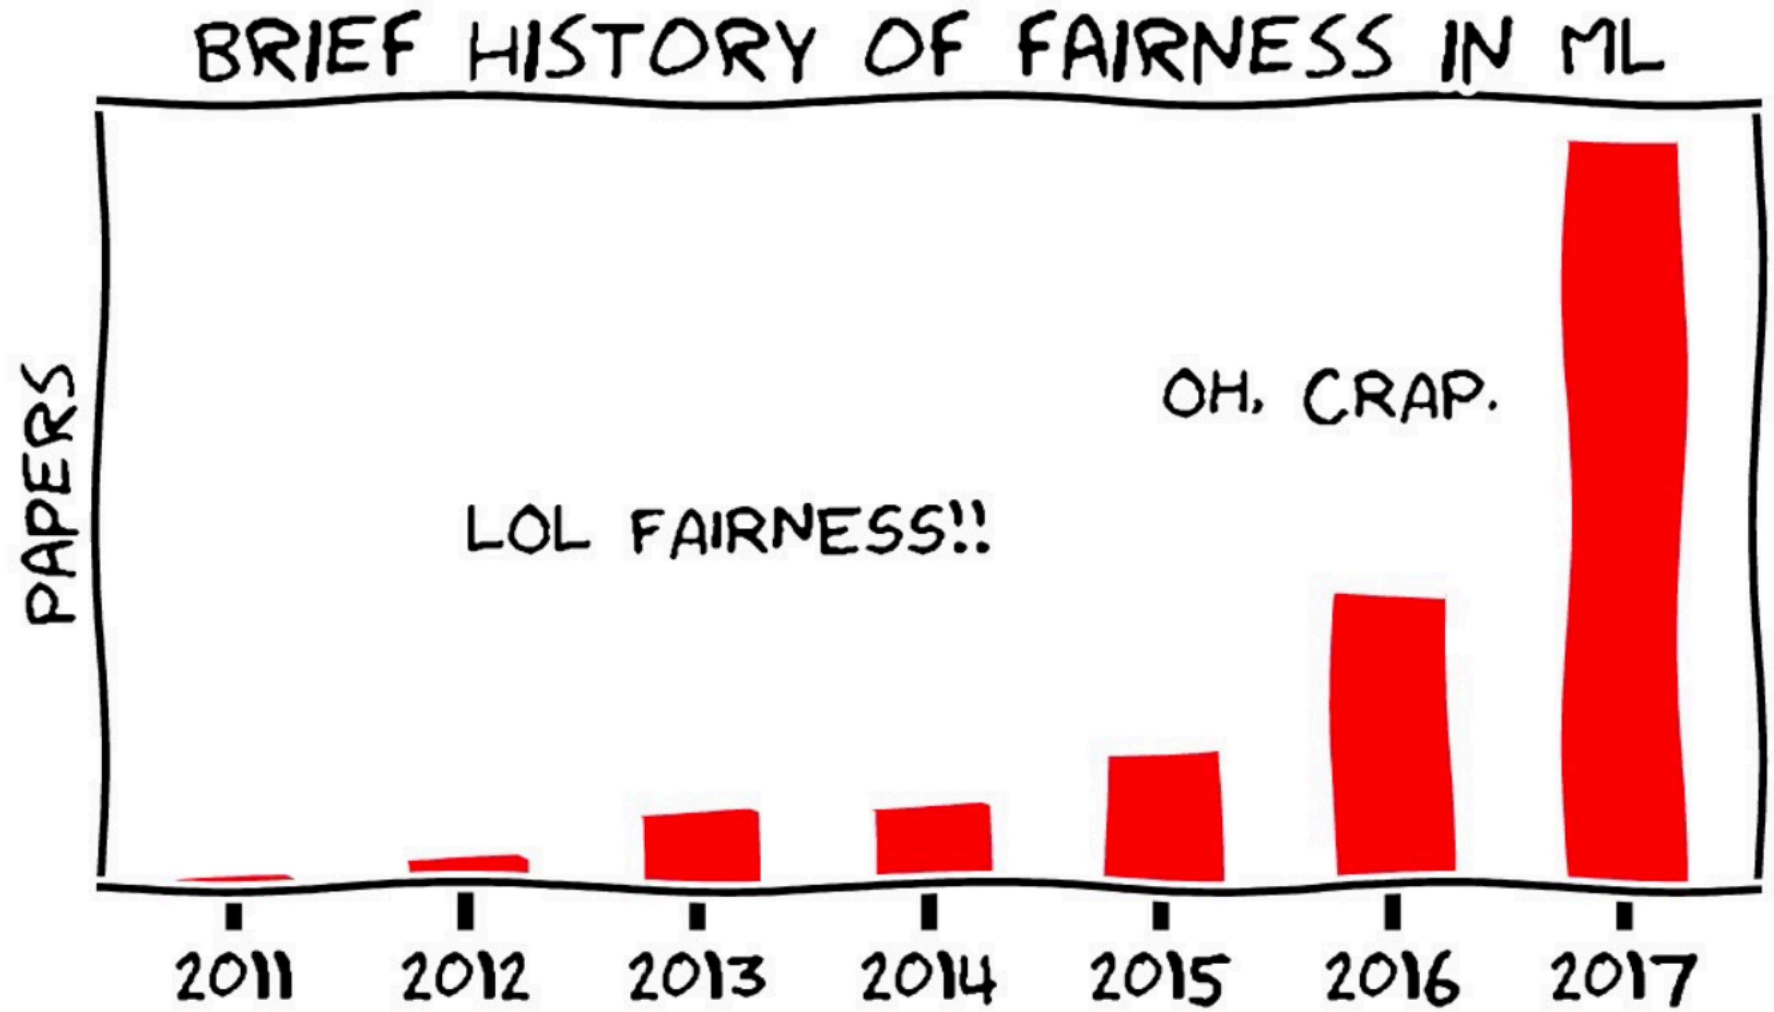

(Credits: Moritz Hardt)

What can go wrong?

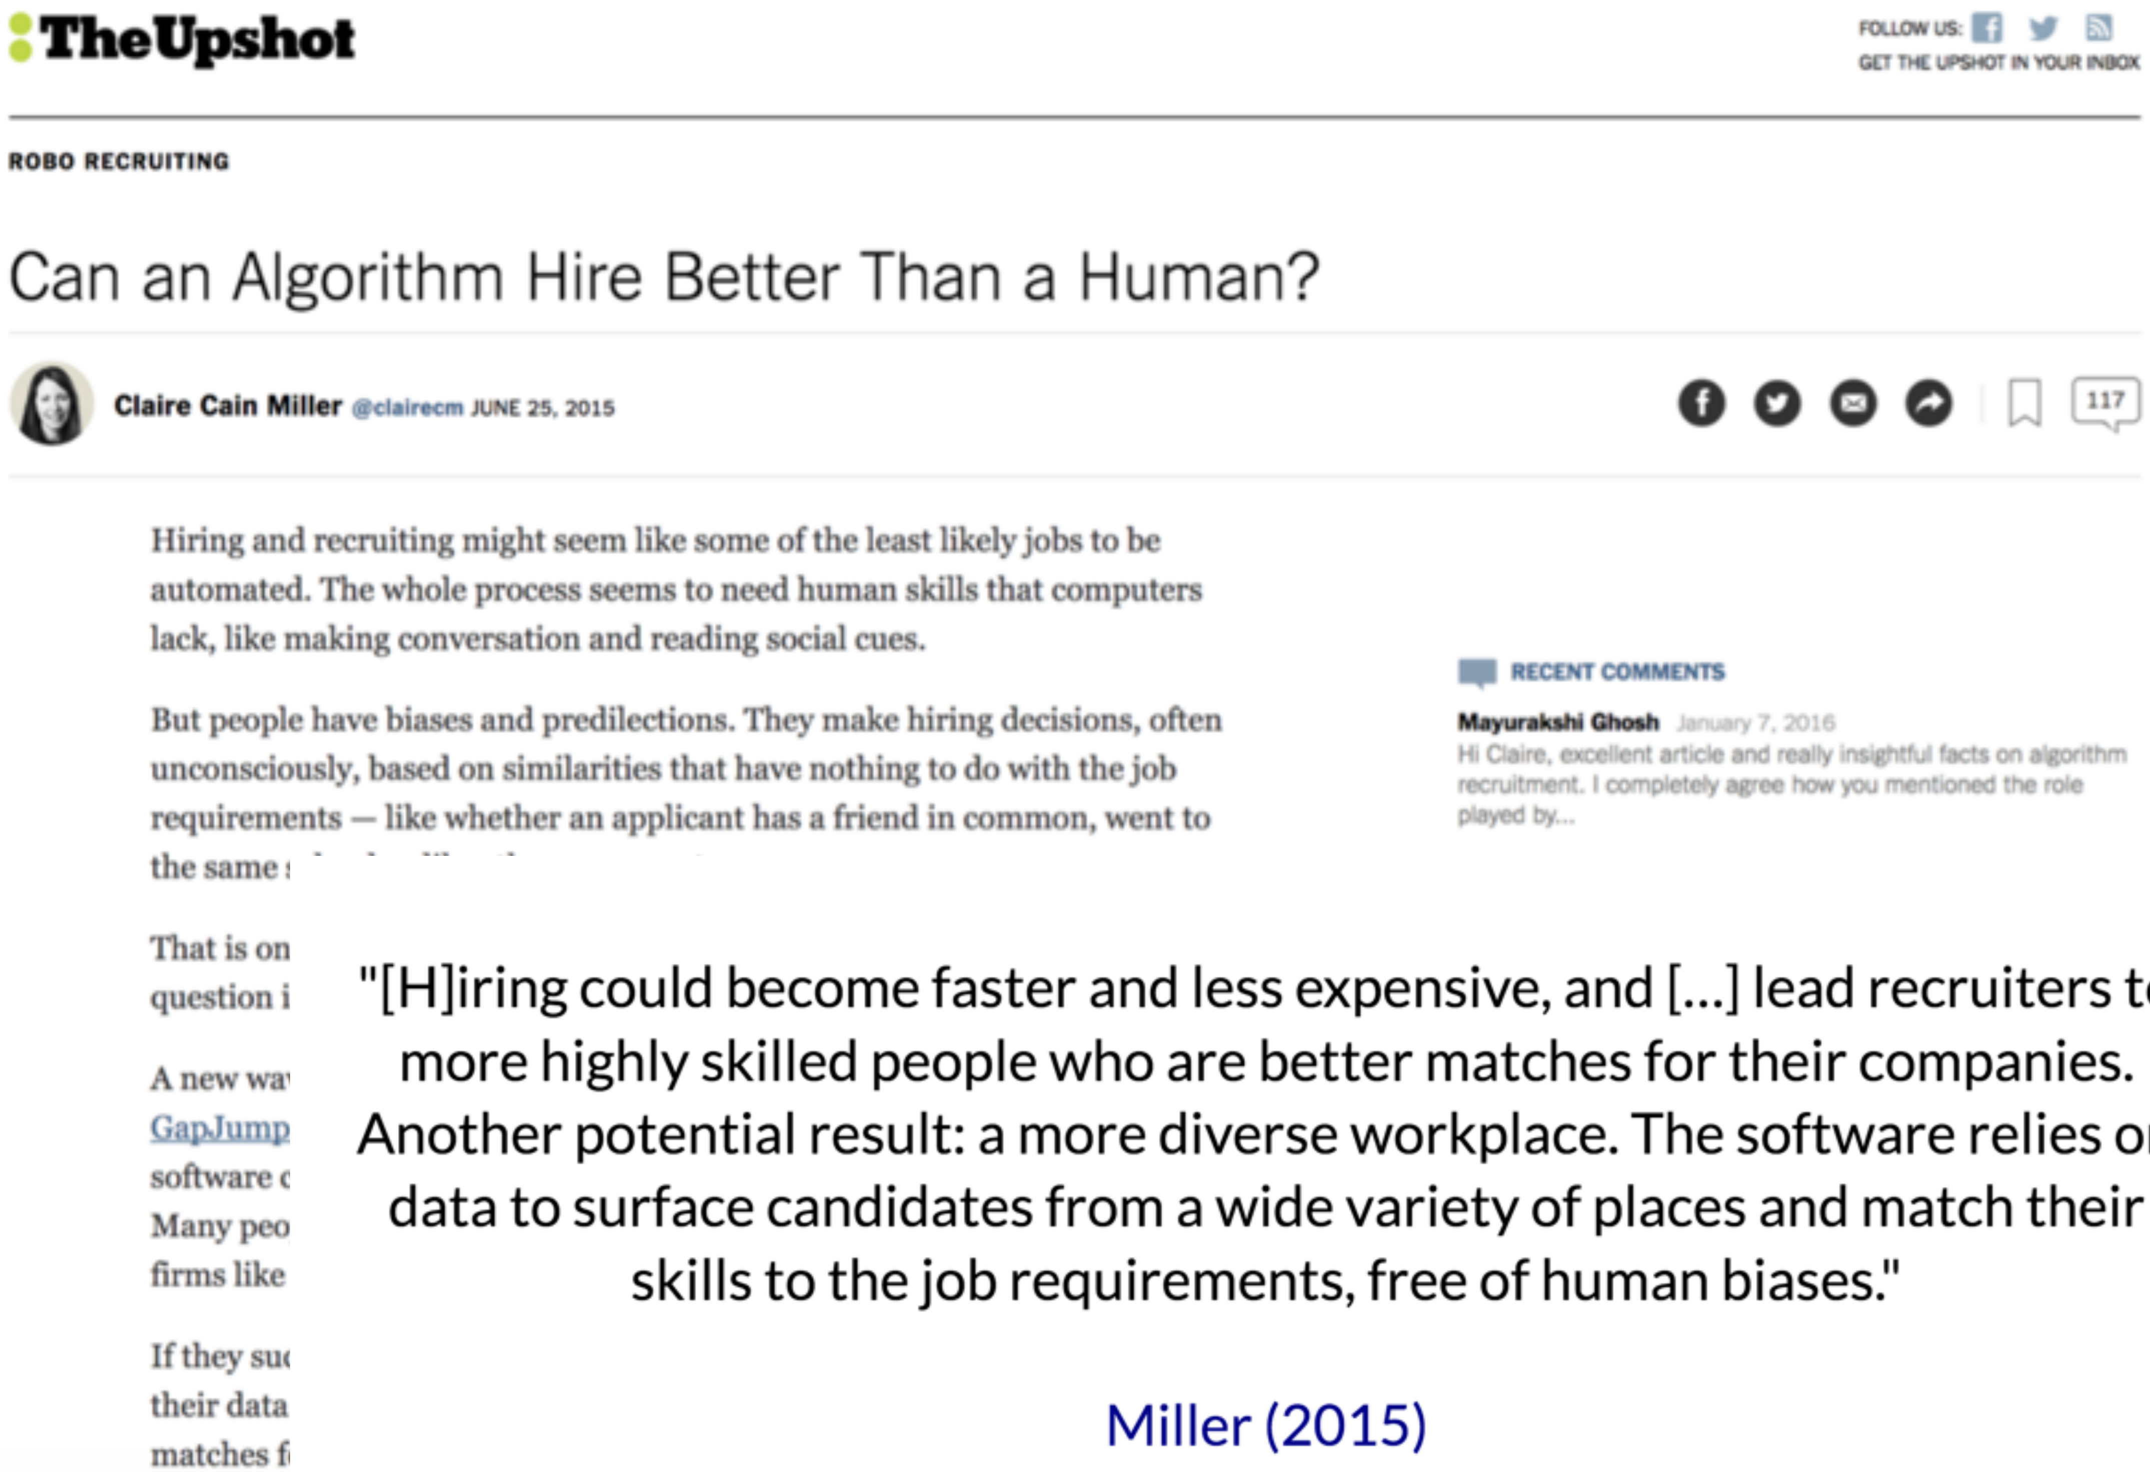

TheUpshot

FOLLOW US:

GET THE UPSHOT IN YOUR INBOX

ROBO RECRUITING

Can an Algorithm Hire Better Than a Human?

Claire Cain Miller @clairecm JUNE 25, 2015

117

Hiring and recruiting might seem like some of the least likely jobs to be automated. The whole process seems to need human skills that computers lack, like making conversation and reading social cues.

But people have biases and predilections. They make hiring decisions, often unconsciously, based on similarities that have nothing to do with the job requirements — like whether an applicant has a friend in common, went to the same s

That is on
question i

A new wa
GapJump
software c
Many peo
firms like

If they suc
their data
matches f

RECENT COMMENTS

Mayurakshi Ghosh January 7, 2016
Hi Claire, excellent article and really insightful facts on algorithm recruitment. I completely agree how you mentioned the role played by...

"[H]iring could become faster and less expensive, and [...] lead recruiters to more highly skilled people who are better matches for their companies. Another potential result: a more diverse workplace. The software relies on data to surface candidates from a wide variety of places and match their skills to the job requirements, free of human biases."

Miller (2015)

TheUpshot

FOLLOW US:
GET THE UPSHOT IN YOUR INBOX

HIDDEN BIAS

When Algorithms Discriminate

Claire Cain Miller @clairecm JULY 9, 2015

147

The online world is shaped by forces beyond our control, determining the

resume screening tool,

the algorithm found two factors to be most indicative of job performance: their name was Jared, and whether they played high school lacrosse.

new study by

Research fro

significantly

names or a h

advertisers a

en and

they do based on people's behavior. As a result [...] algorithms can reinforce human prejudices."

Miller (2015)

(Credits: Moritz Hardt)

Search query	Work experience	Education experience	Profile views	Candidate	Xing ranking
Brand Strategist	146	57	12992	male	1
Brand Strategist	327	0	4715	female	2
Brand Strategist	502	74	6978	male	3
Brand Strategist	444	56	1504	female	4
Brand Strategist	139	25	63	male	5
Brand Strategist	110	65	3479	female	6
Brand Strategist	12	73	846	male	7
Brand Strategist	99	41	3019	male	8
Brand Strategist	42	51	1359	female	9
Brand Strategist	220	102	17186	female	10

TABLE II: Top k results on www.xing.com (Jan 2017) for the job serach query "Brand Strategist".

	WHITE	AFRICAN AMERICAN
Labeled Higher Risk, But Didn't Re-Offend	23.5%	44.9%
Labeled Lower Risk, Yet Did Re-Offend	47.7%	28.0%

Overall, Northpointe's assessment tool correctly predicts recidivism 61 percent of the time. But blacks are almost twice as likely as whites to be labeled a higher risk but not actually re-offend. It makes the opposite mistake among whites: They are much more likely than blacks to be labeled lower risk but go on to commit other crimes. (Source: ProPublica analysis of data from Broward County, Fla.)

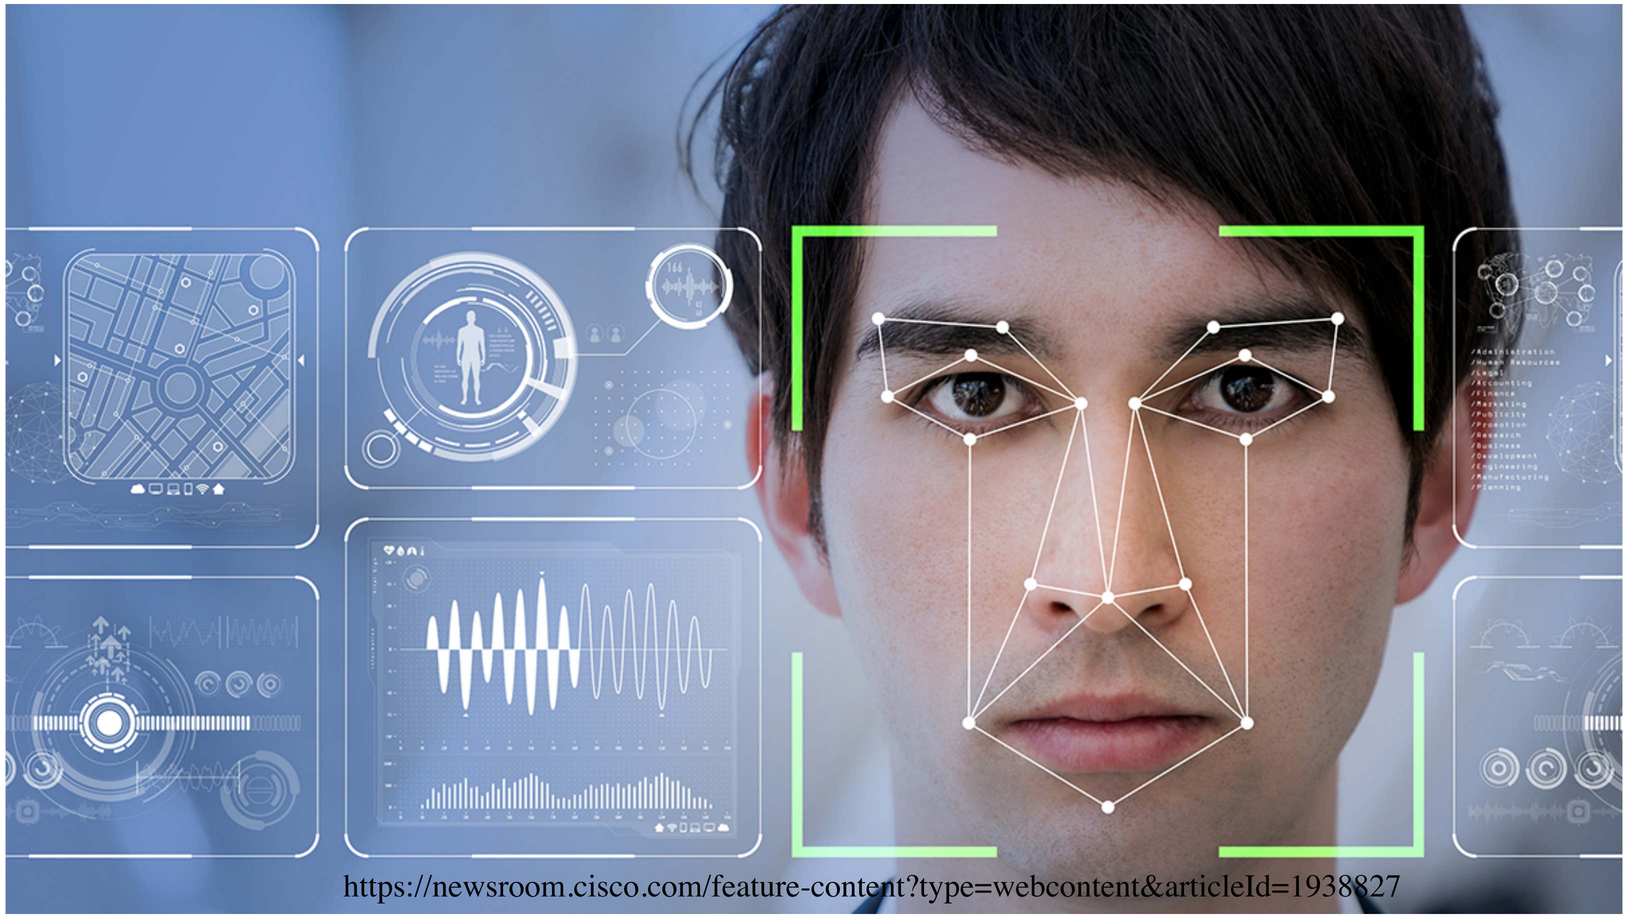

The New York Times

Facial Recognition Is Accurate, if You're a White Guy

By Steve Lohr

Gender Classifier	Darker Male	Darker Female	Lighter Male	Lighter Female	Largest Gap
Microsoft	94.0%	79.2%	100%	98.3%	20.8%
FACE++	99.3%	65.5%	99.2%	94.0%	33.8%
IBM	88.0%	65.3%	99.7%	92.9%	34.4%

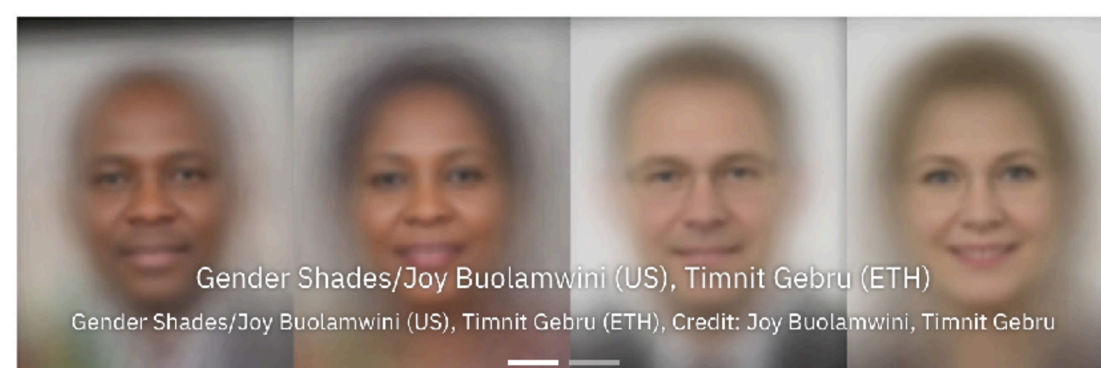

Why is this happening?

→ Skewed sample

→ Tainted examples

→

Man is to Computer Programmer as Woman is to Homemaker?
Debiasing Word Embeddings

Tolga Bolukbasi[1], Kai-Wei Chang[2], James Zou[2], Venkatesh Saligrama[1,2], Adam Kalai[2]
[1]Boston University, 8 Saint Mary's Street, Boston, MA
[2]Microsoft Research New England, 1 Memorial Drive, Cambridge, MA
tolgab@bu.edu, kw@kwchang.net, jamesyzou@gmail.com, srv@bu.edu, adam.kalai@microsoft.com

→ Sample size disparity

[Barocas Selbst 2016]

How to fix this?

Natural solution:

Fairness Through Blindness

"We don't even look at 'race'!"

Problem: You do not need to look at 'race' to predict the race

→ Proxies are everywhere
(and ML is great at picking them up)

It's Not Privacy, and It's Not Fair

Cynthia Dwork & Deirdre K. Mulligan *

What is fair?

Legally recognized 'protected classes'

Race (Civil Rights Act of 1964); **Color** (Civil Rights Act of 1964); **Sex** (Equal Pay Act of 1963; Civil Rights Act of 1964); **Religion** (Civil Rights Act of 1964); **National origin** (Civil Rights Act of 1964); **Citizenship** (Immigration Reform and Control Act); **Age** (Age Discrimination in Employment Act of 1967); **Pregnancy** (Pregnancy Discrimination Act); **Familial status** (Civil Rights Act of 1968); **Disability status** (Rehabilitation Act of 1973; Americans with Disabilities Act of 1990); **Veteran status** (Vietnam Era Veterans' Readjustment Assistance Act of 1974; Uniformed Services Employment and Reemployment Rights Act); **Genetic information** (Genetic Information Nondiscrimination Act)

https://mrtz.org/nips17/#/10

→ Disparate treatment

→ Disparate outcome

But: How to make that precise?

→ Unawareness

→ Demographic Parity

→ Equalized Odds

→ Predictive Rate Parity

→ Individual Fairness

→ Counterfactual Fairness

(Big) Problem:

Fundamental incompatibility

Is it only about fairness?

https://www.youtube.com/watch?v=ecClODh4zYk

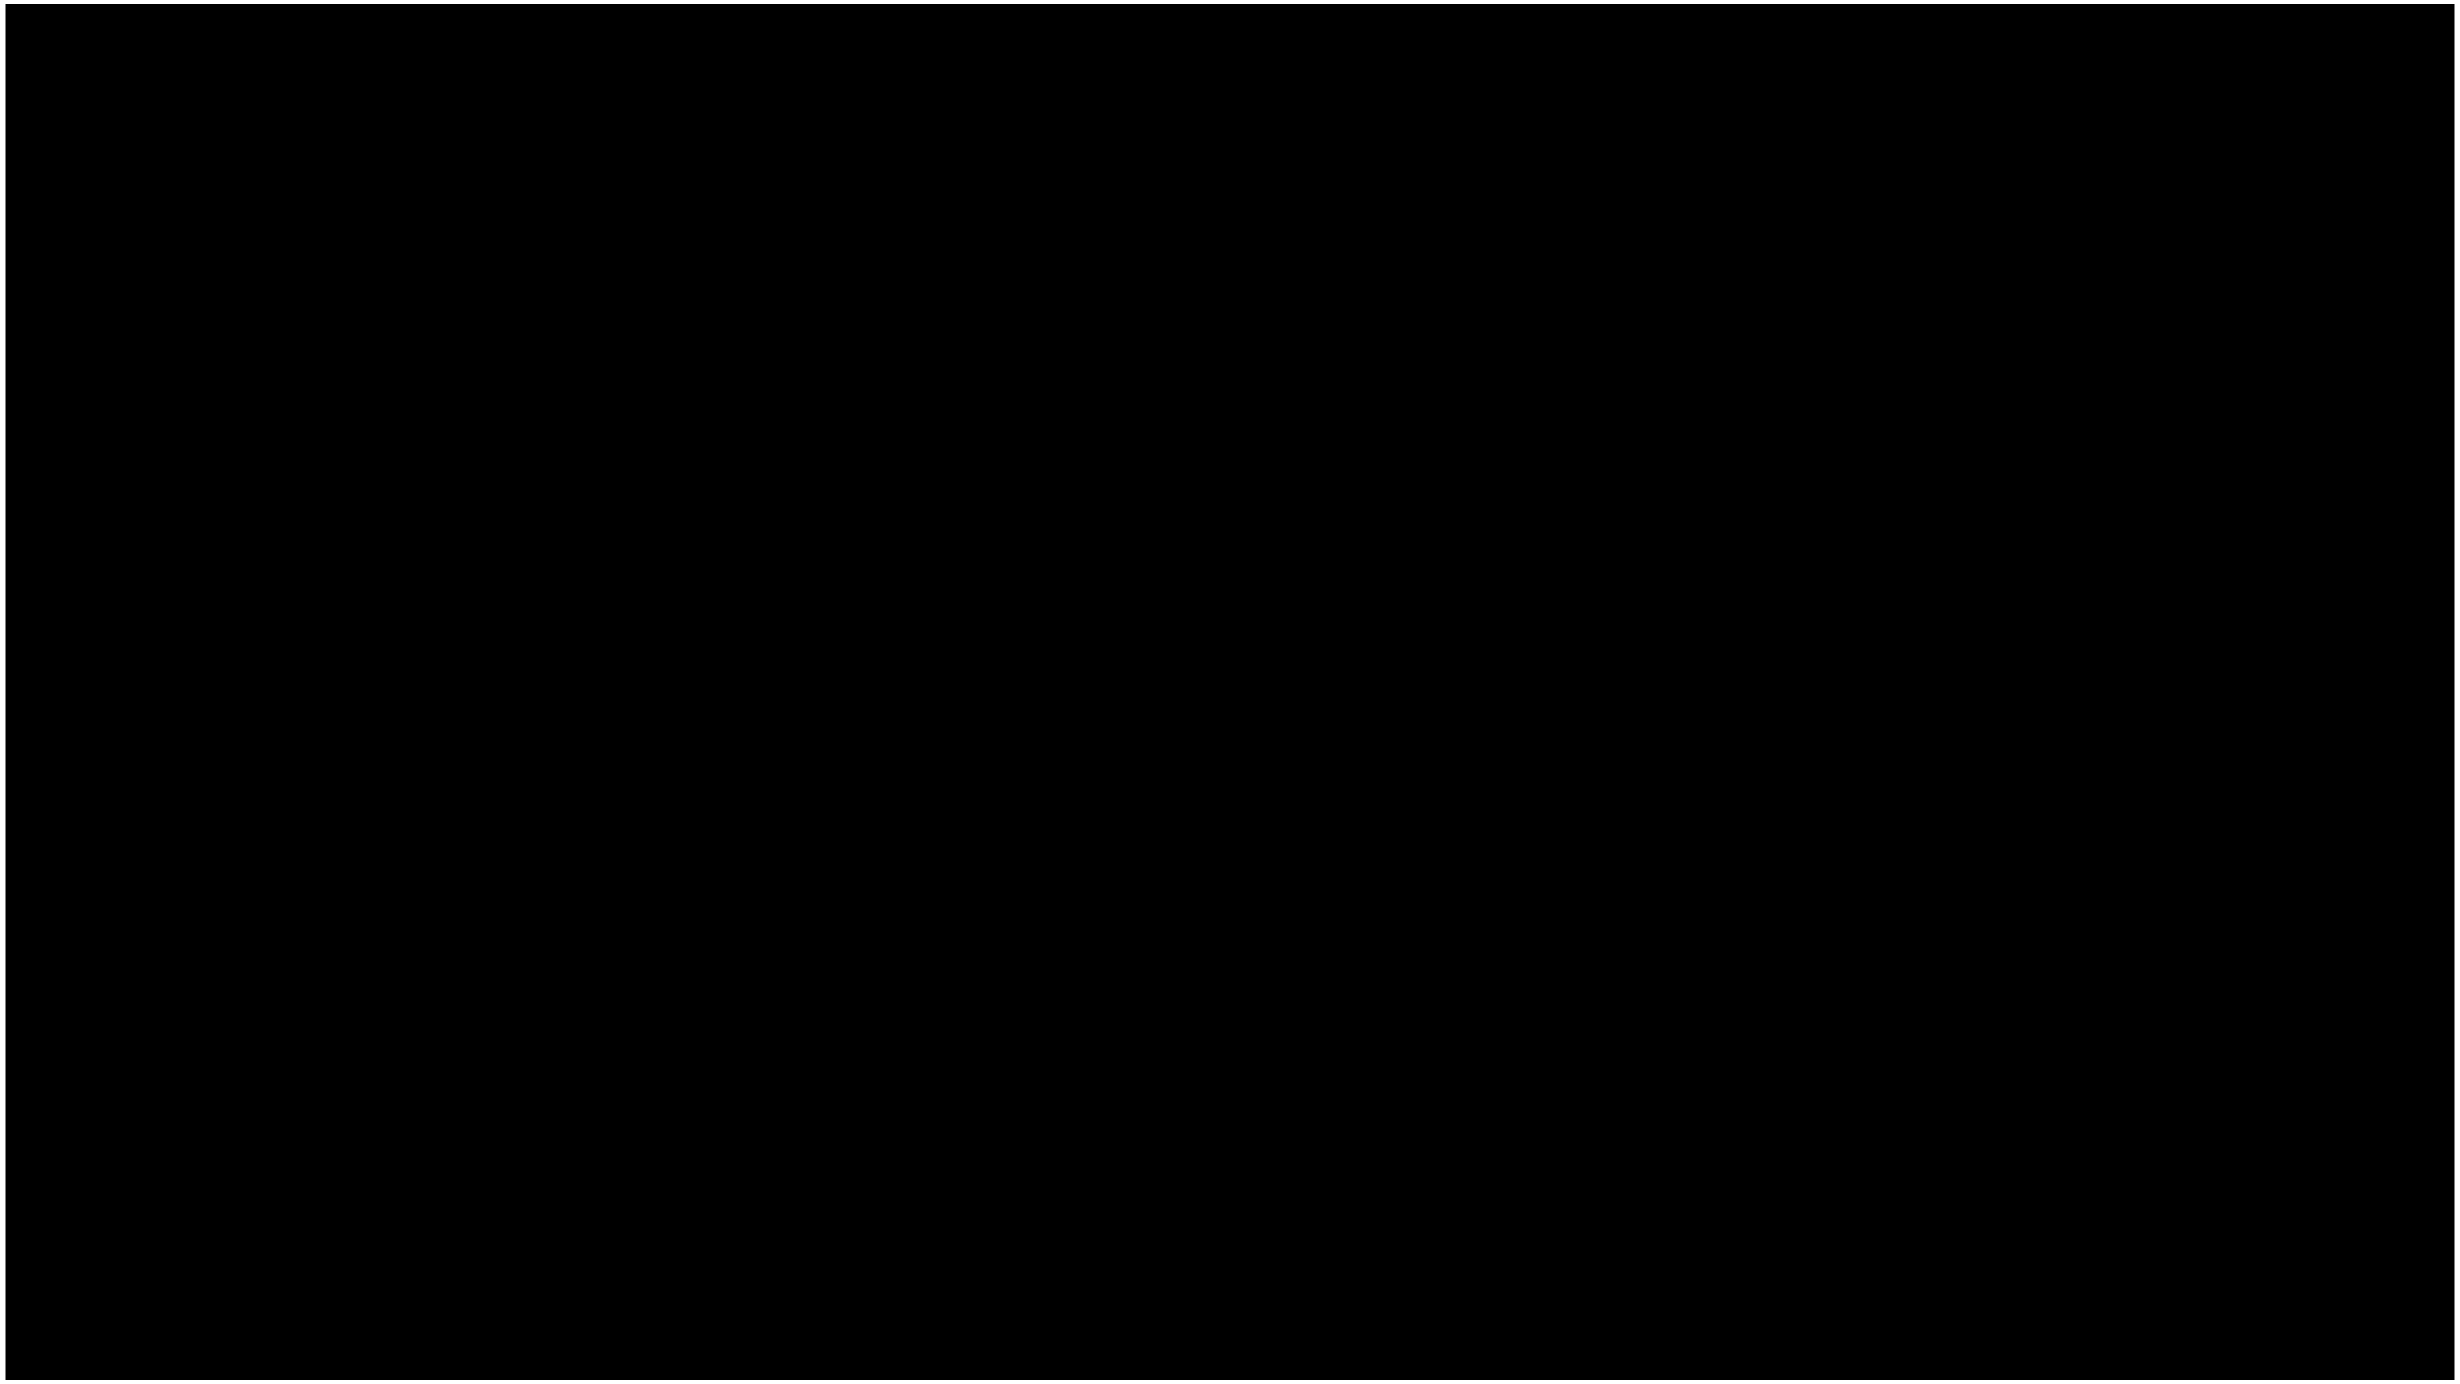

https://www.youtube.com/watch?v=bWOUIvknFBc