



How digital technologies are changing work

Why the world needs to worry about the ethics of AI



SPEAKER:

Simon Greenman

Member, Global AI Council, World Economic Forum

Partner, Best Practice AI

Follow us on Twitter:

@digitcentre #DigitDebates



**Sign up for our newsletter at
digit-research.org**

AI is a GPT that will be woven into the fabric of society with the potential to transform lives, companies, and government


Improving crop yields in India with predictive plant disease diagnosis



Predicting occurrences of diseases earlier and more accurately



Improving corporate performance – CV screening & supply chain forecasting



The total global uplift on GDP as a result of AI could be over \$10 trillion by 2030 per PWC

North America

\$3.7 Trillion
15% of GDP

North Europe

\$1.8 Trillion
10% of GDP

China

\$7.0 Trillion
26% of GDP

Overcoming Racial Bias In AI Systems And Startlingly Even In AI Self-Driving Cars

Racial bias in a medical algorithm favors white patients over sicker black patients

AI expert calls for end to UK use of 'racially biased' algorithms

AI Bias Could Put Women's Lives At Risk - A Challenge For Regulators

Gender bias in AI: building fairer algorithms

Bias in AI: A problem recognized but still unresolved

Amazon, Apple, Google, IBM, and Microsoft worse at transcribing black people's voices than white people's with AI voice recognition, study finds

Millions of black people affected by racial bias in health-care algorithms

Study reveals rampant racism in decision-making software used by US hospitals – and highlights ways to correct it.

When It Comes to Gorillas, Google Photos Remains Blind

Google promised a fix after its photo-categorization software labeled black people as gorillas in 2015. More than two years later, it hasn't found one.

Google 'fixed' its racist algorithm by removing gorillas from its image-labeling tech

The Week in Tech: Algorithmic Bias Is Bad. Uncovering It Is Good.

Artificial Intelligence has a gender bias problem – just ask Siri

The Best Algorithms Struggle to Recognize Black Faces Equally

US government tests find even top-performing facial recognition systems misidentify blacks at rates five to 10 times higher than they do whites.

WILL A ROBOT STEAL YOUR JOB?

KILLER COMPUTERS

Bill Gates warns 'dangerous AI' poses a threat 'like nuclear weapons'

AI WARNING:
Robots will destroy a HUGE number of jobs, claims expert


AI could be used to **TAKE OVER** the **WORLD** through 'evil' fake news and hijacking cars

<https://phvs.org/news/2019-10-headlines-dont-robots-threaten-jobs.html>

**We have a lot of fear
and not enough trust
of AI. Why?**



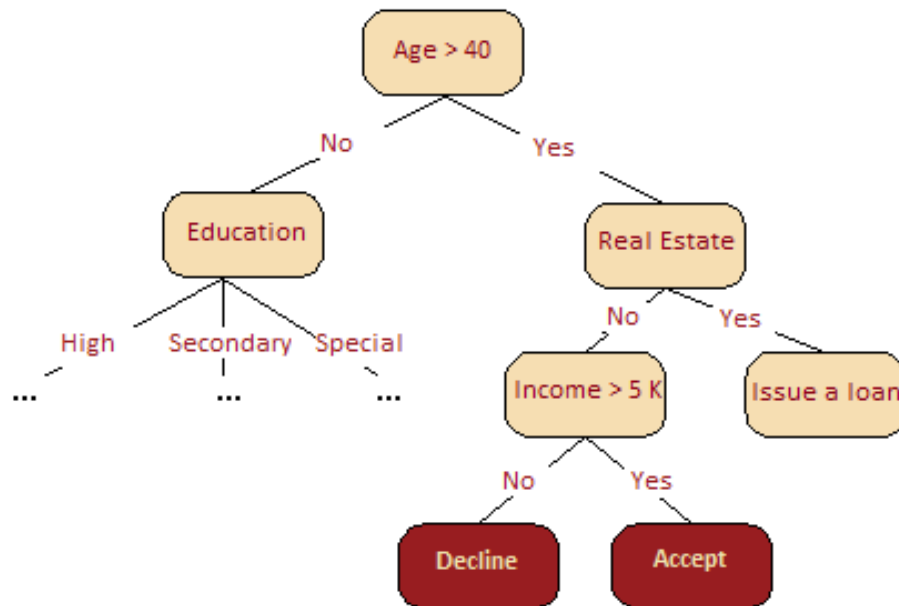
AI and Machine Learning creates its own “software rules” by learning patterns in data




And we do not necessarily understand the “rules” it creates

There is often a tradeoff between interpretability and accuracy of explainability

Inferable models, such as decision trees, provide non-technical explanations (e.g. a loan decision)



Deep neural networks are often described as black boxes and it is harder to understand their decision making



How do you understand a 175 billion parameter deep neural network model? OpenAI's GPT-3 terrified us in 2020 as it wrote “intelligently”

OpenAI 175 billion parameters (June '20)

A robot wrote this entire article. Are you scared yet, human?

We asked GPT-3, OpenAI's powerful new language generator, to write an essay for us from scratch. The assignment? To convince us robots come in peace

- For more about GPT-3 and how this essay was written and edited, please read our editor's note below



▲ 'We are not plotting to take over the human populace.' Photograph: Volker Schlichting/Getty Images/EyeEm

Google 1.6 trillion parameters (Jan '21)

VB The Machine GamesBeat Jobs Special Issue Become a Member | Sign In

The Machine Making sense of AI

Google trained a trillion-parameter AI language model

Kyle Wiggers @Kyle_L_Wiggers January 12, 2021 10:36 AM

f t in

NEW YORK, NEW YORK - OCTOBER 20: Google's offices stand in downtown Manhattan on October 20, 2020 in New York City. Accusing the company of using anticompetitive tactics to illegally monopolize the online search and search advertising markets, the Justice Department and 11 states Tuesday filed an antitrust case against Google. Image Credit: Spencer Platt/Getty Images

How open banking is driving huge innovation

Learn how fintechs and forward-thinking FIs are accelerating personalized financial products through data-rich APIs.

Parameters are the key to machine learning algorithms. They're the part of the model that's learned from historical training data. Generally speaking, in the language domain, the correlation between the number of parameters and sophistication has held up remarkably well. For example, OpenAI's GPT-3 — one of the largest language models ever trained, at 175 billion parameters — can make primitive analogies, generate recipes, and even complete basic code.

A futuristic robot with a white, human-like face and a chrome, mechanical body. The robot is shown in profile, holding its head with one hand, suggesting a state of distress, pain, or deep thought. The background is dark, and the text "How can we trust AI when it fails?" is overlaid in white, bold, sans-serif font.

How can we trust AI when it fails?

Deep learning is brittle and lacks human level robustness

It recognises statistical patterns, not higher order concepts and lacks common sense



School Bus

100%



Garbage
Truck

99%



Punch Bag

100%



Snow Plough

92%

Tesla's Smart Summons shows the brittleness of pattern recognition and the challenges of a world of edge cases



And when it fails it goes viral on social media

我们 WEVIDEO


微博
@紧急呼叫

12月25日 福州中防万宝城商场
导购机器人跌落扶梯 撞翻乘客



**Is AI
biased?**

Joy Buolamwini is on a mission to “stop an unseen force that is rising.” The risk of bias and discrimination.



Gender classification systems are often biased as they are not trained on representative sample datasets of gender & ethnicity

Gender was misidentified in up to 1% of lighter-skinned males



Gender was misidentified in up to 7% of lighter-skinned females



Gender was misidentified in 35% of darker skinned females




These biases could have serious ethical, legal, operational and reputational consequences

The Atlantic

AI-Driven Dermatology Could Leave Dark-Skinned Patients Behind

Machine learning has the potential to save thousands of people from skin cancer each year—while putting others at greater risk.

ANGELA LASHBROOK AUGUST 16, 2018



STEVE GSCHMEISSNER / GETTY

LaToya Smith was 29 years old when she died from skin cancer. The young doctor had gotten her degree in podiatry from Rosalind Franklin University, in Chicago, just four

AI shines a spotlight and often amplifies our human biases

What did the algorithm learn after reading 3.5 million books and 11 billion words?

Top 11 positive and negative words associated with females and males

Female	
Positive	Negative
beautiful	battered
lovely	untreated
chaste	barren
gorgeous	shrewish
fertile	sheltered
beauteous	heartbroken
sexy	unmarried
classy	undernourished
exquisite	underweight
vivacious	uncomplaining
vibrant	nagging

Male	
Positive	Negative
just	unsuitable
sound	unreliable
righteous	lawless
rational	inseparable
peaceable	brutish
prodigious	idle
brave	unarmed
paramount	wounded
reliable	bigoted
sinless	unjust
honorable	brutal

“There is a systemic, systematic, racist, sexist, gendered, class-oriented and other axes of discrimination-bias embedded in most data collected by humans”

<https://news.sky.com/story/top-civil-servant-at-department-for-education-to-leave-after-a-levels-chaos-12057141>

Government leaders
politicising algorithms
doesn't help to engender
public trust in AI!

sky news

Boris Johnson blames 'mutant
algorithm' for A-level results fiasco, then
sacks education chief

The prime minister is accused of "shamelessly trying to avoid taking responsibility"
for the exams debacle.



**Is our AI
technology
human centric?**



A young man with dark hair is looking down at a smartphone he is holding. The scene is dimly lit with a strong red light source from the right, casting a glow on his face and shirt. Overlaid on the image are several semi-transparent social media-style elements. On the left, there's a post with three photos of young women, a caption 'I love these two!', and a heart count of 7. On the right, there's a post with a photo of a woman, a caption 'I can't wait for this weekend!', and a heart count of 3. The background is dark and out of focus.

THE TECHNOLOGY THAT CONNECTS US ALSO CONTROLS US

**/the
social
dilemma_**

| NETFLIX

[thesocialdilemma.com](https://www.thesocialdilemma.com)

Do we understand the (unintended) consequences of our algorithms?



200+ “Ethical AI” frameworks from the likes of the OECD and the IEEE boil down to similar principles

- 1) **Explainable** and transparent decision making
- 2) **Inclusive**, diverse and **fair** (avoid or don't reinforce bias)
- 3) Be built and tested for **safety**
- 4) Be **socially** beneficial
- 5) **Respect** human rights and the law
- 6) People are **accountable**

AI regulation is already here under GDPR: Explainability

Article 22 under GDPR states:

Fully automated decisions with legal effect or similarly significant effect needs to be explainable

and

data subjects have the right to human-made decisions




Explainability will spotlight company historical practices

Information Commissioners Office (ICO) – requires an Explainability statement for automation with legal effect

The ICO identified six main types of explanation that form an explanatory statement:

1. **Rationale** explanation: the reasons that led to a decision, delivered in an accessible and non-technical way.
2. **Responsibility** explanation: who is involved in the development, management and implementation of an AI system, and who to contact for a human review of a decision.
3. **Data** explanation: what data has been used in a particular decision and how; what data has been used to train and test the AI model and how.
4. **Fairness** explanation: steps taken across the design and implementation of an AI system to ensure that the decisions it supports are generally unbiased and fair, and whether or not an individual has been treated equitably.
5. **Safety** and performance explanation: steps taken across the design and implementation of an AI system to maximise the accuracy, reliability, security and robustness of its decisions and behaviours.
6. **Impact** explanation: the impact that the use of an AI system and its decisions has or may have on an individual, and on wider society.



Companies are implementing Responsible AI programmes, often as a competitive differentiator

- **Deutsche Telekom view digital ethics as a strategic competitive differentiator**
- **They focus on shaping AI responsibly, with programme implemented on various levels**
- **Continuously updated**



INTERNAL PROCESSES

Integration into internal security and data protection processes; integration into financing processes



DIGITAL ETHICS CENTER

Place for internal/external conferences to make Digital Ethics tangible



DIGITAL ETHICS SEAL

First internal AI projects have been certified and received a Digital Ethics Seal



SUPPLIER MANAGEMENT

Guidelines extended to suppliers of AI systems



COMMUNICATION & EVENTS

Regular communication and expert interviews on the intranet; AI days and other formats




TRAININGS & FORMAL POLICY

eLearnings, roadshow, workshops and policy for employees to develop safe AI systems



Can the identification of bias by AI help us create a more just society?

- A recent paper looked at disparities in the treatment of knee osteoarthritis that causes chronic pain
- Radiologists review x-rays of the knee and score the patient's pain based on radiographic features (e.g. degree of missing cartilage or structural damage)
- But there is a gap between the radiologist prediction of pain and self-reporting pain of black patients
- Why? The methodology used - the Kellgren-Lawrence (KLG) grade - was developed several decades ago based on a white British population
- ML used to predict pain in black patients was much more accurate – medical methods need to be updated.




**BIG BROTHER
IS WATCHING
YOU**



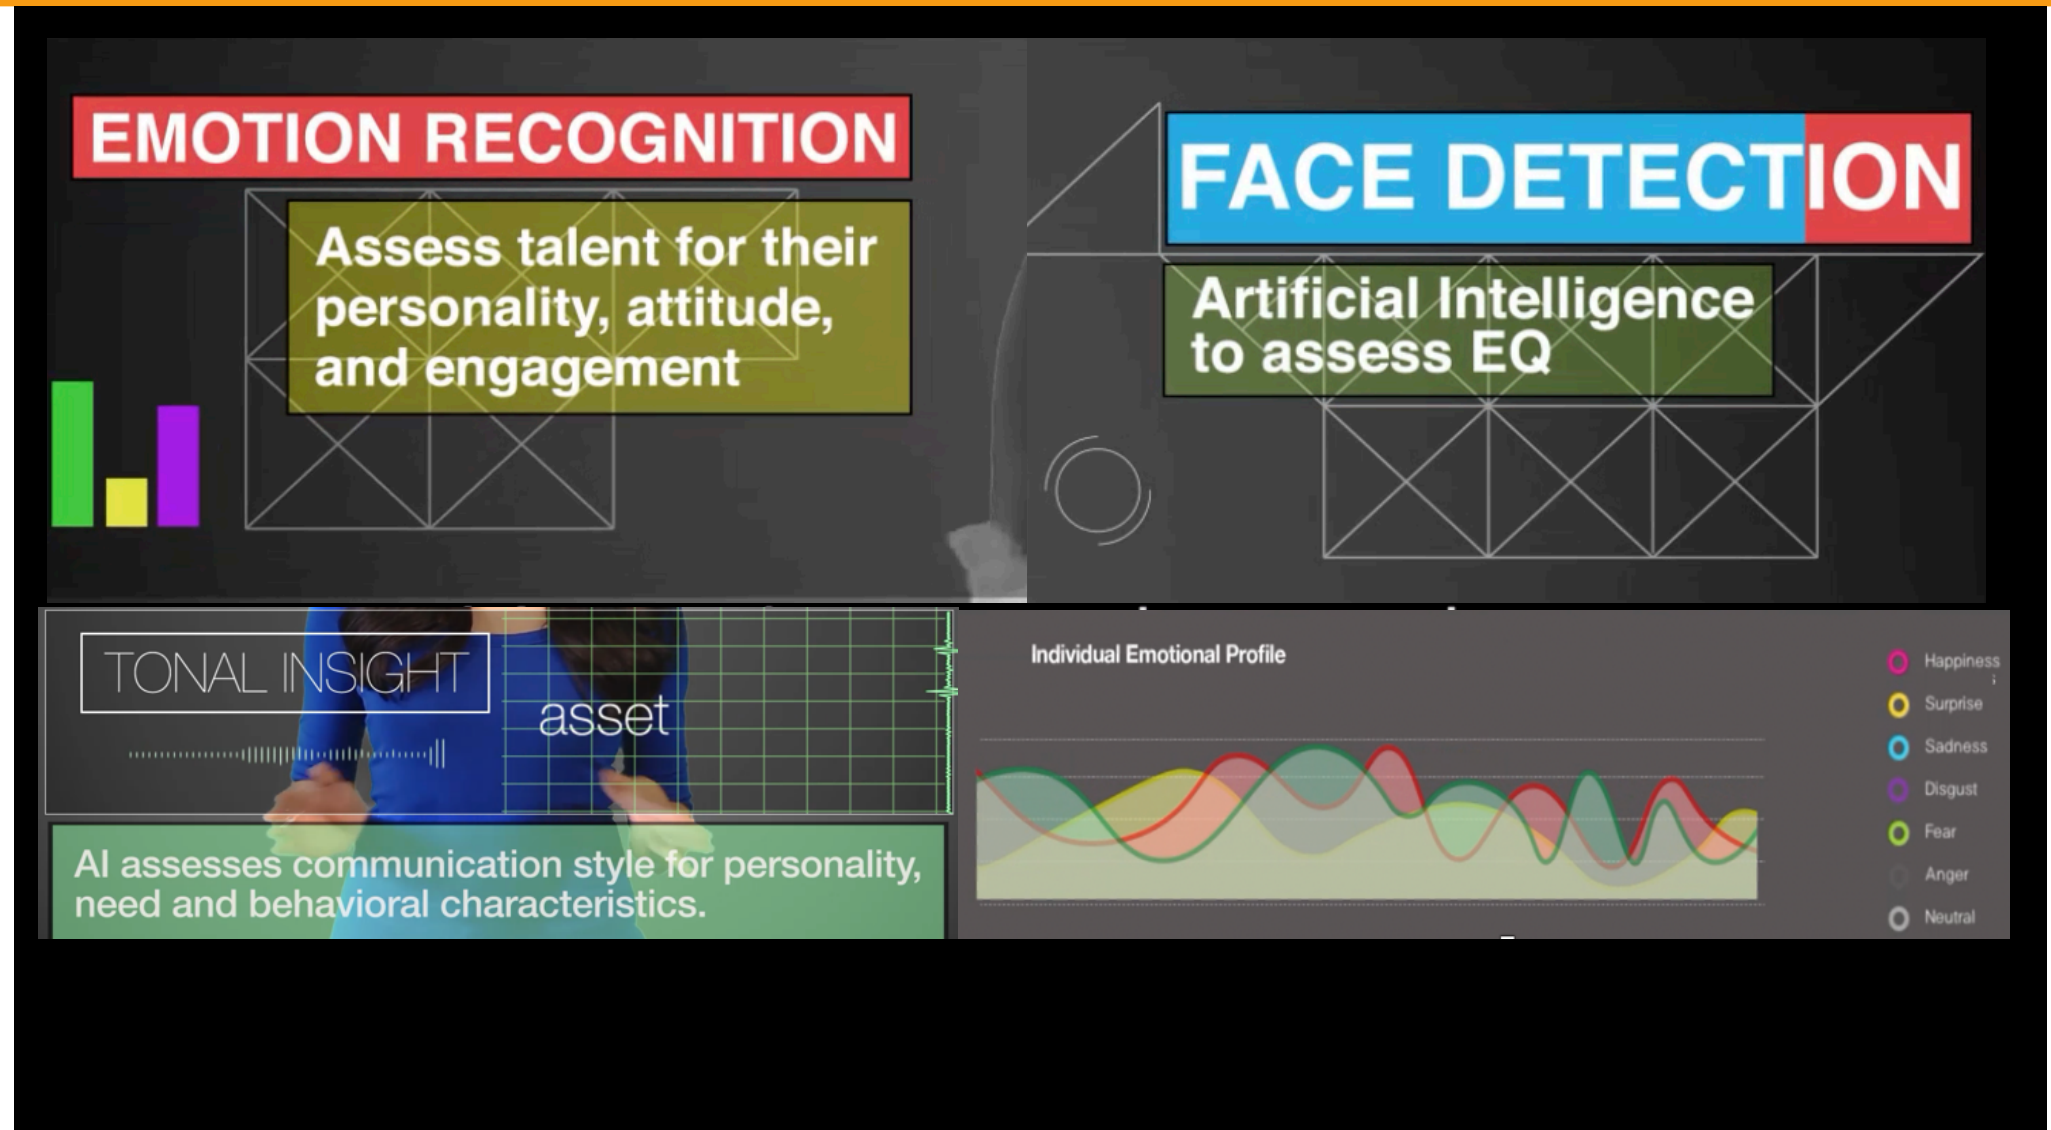
OBEY

**What does
society
deem
acceptable
use cases
for AI? ...**

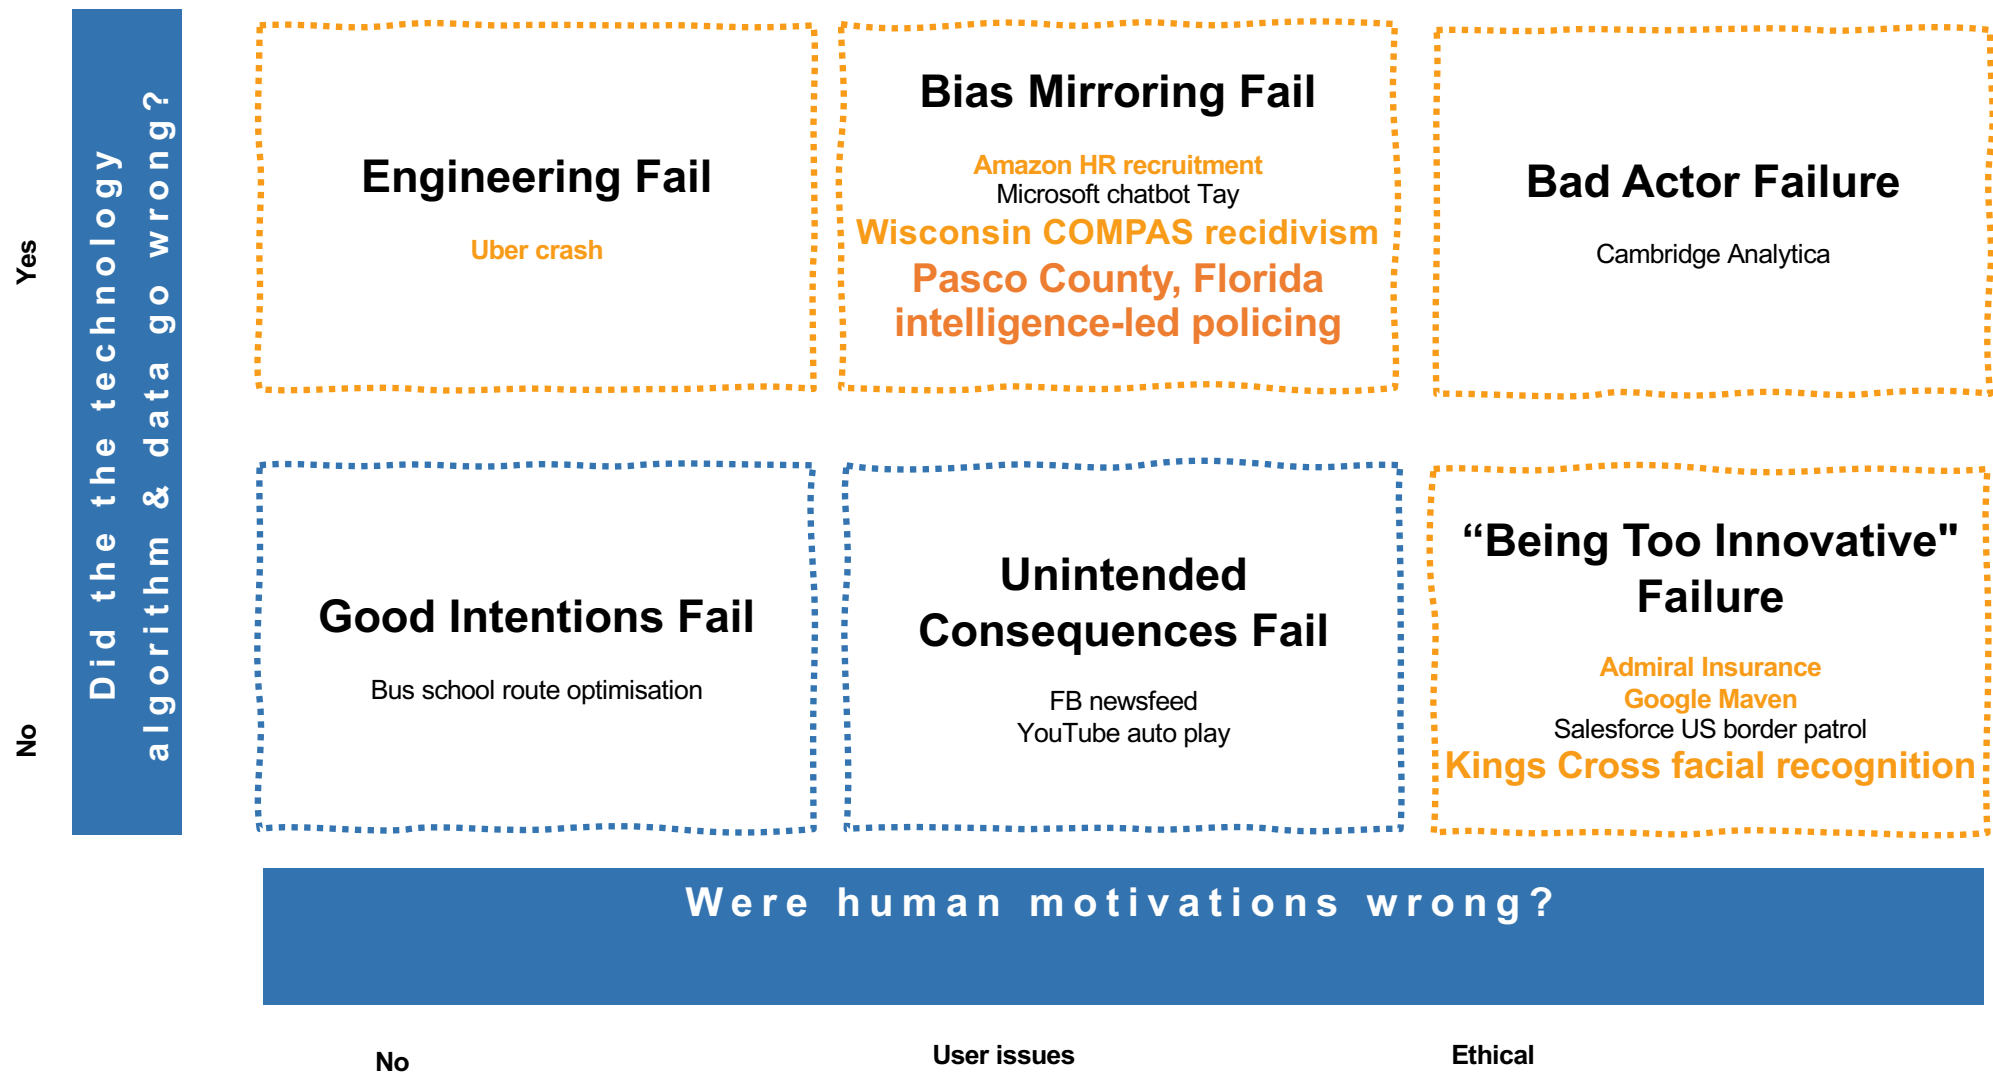


What about algorithmic monitoring and assessment?


How about we save time in recruiting with automated video interviews?



AI ethics is often more a question of human motivations than technology failings



AI ethics is often more a question of human motivations than technology failings

<u>Did the technology, algorithm or data go wrong?</u>	Yes	No
		
	No	Yes
	Were human motivations wrong?	

AI surveillance helped many countries successfully manage COVID-19. Should we use AI monitoring in our country?

Thermal screening in stations and airports




FRT to detect those wearing masks & enforce quarantines



Thermal cameras used on on drones to identify fevers and crowds



Enforcing social distancing



What if we could understand what everyone was saying in our company? Should we?

What if....

- **All the conversations taking place in your organisation were tracked**
 - Emails, slack, social media, telephone, messaging, meetings...
 - Customers, teams, suppliers, feedback sessions, internal debate....
- **What value could your organisation capture if these were processed and assessed**
 - Pain points and customer concerns identified, issues logged, ranked and prioritised...
 - Sentiment analysed, management issues spotted, individual performance assessed...
 - Personalised feedback, abuse monitoring, behaviour nudging, welfare protected...
- **But at what cost?**
 - Privacy, nuance, personal style...
 - Transparency versus reality...
- **Who has the power – and where might this happen?**

But how about AI in education? Should we monitor students in the classroom?

- AI can help education. For example, **VIPKid** is used by 700,000+ students. Its AI animated "fun characters" assist human teachers. Correct answers grew 50 to 80%.
- Chinese **Megvii** is a \$4B+ company known for its facial recognition platform Face++
- Demonstrated a **classroom teaching evaluation system** used to supplement teaching evaluations through **real-time structured analysis** of classroom video data
- Observes and **classifies** student behaviour:
 - resting on one's desk
 - playing on a mobile phone
 - sleeping
 - listening to a lecture
 - reading
 - raising hands
 - concentration levels...

A recent report identified 10+ Chinese companies in the emotion recognition market for education

EF Children's English
英孚少儿英语

In person and
online

Partners with Tencent Cloud to conduct image, emotion, and voice recognition, and receives curriculum design assistance to EF's product-development teams and teachers.¹⁴¹

Hanwang Education
汉王教育

In person

Class Care System (CCS) cameras take photos of whole classes once per second, connect to a programme that purportedly uses deep-learning algorithms to detect behaviours (including 'listening, answering questions, writing, interacting with other students, or sleeping') and issue behavioural scores to students every week. Scores are part of

New Oriental
新东方

Blended learning

AI Dual Teacher Classrooms contain a 'smart eye system based on emotion recognition and students' attention levels', which the company says can also detect emotional states, including 'happy, sad, surprised, normal, and angry'.¹⁴⁹ A

Hikvision
海康威视

In person

Smart Classroom Behaviour Management System integrates three cameras, positioned at the front of the classroom, and identifies seven types of emotions (fear, happiness, disgust, sadness, surprise, anger, and neutral) and six behaviours (reading, writing, listening, standing, raising hands, and laying one's head on a desk).¹⁴⁴ Cameras take attendance using face

Meezao
蜜枣网

In person

Uses facial expression recognition and eye-tracking software to scan preschoolers' faces over 1,000 times per day and generate reports, which are shared with teachers and parents.¹⁴⁷ Reports contain data visualisations of students'

Taigusys Computing
太古计算

In person

Collects data from three cameras, one each on students' faces, teachers, and a classroom's blackboard. The system detects seven emotions (neutral, happy, surprised, disgusted, sad, angry, scared) and seven actions (reading, writing, listening, raising hands, standing up, lying on the desk, playing with mobile phones).¹⁵¹

<https://www.article19.org/wp-content/uploads/2021/01/ER-Tech-China-Report.pdf>

And many Chinese vendors are offering emotion recognition and monitoring for "public security." Any concerns?


Alpha Hawkeye 阿尔法鹰眼	Monitors vestibular emotional reflex and conducts posture, speech, physiological, and semantic analysis. ⁶⁵	<ul style="list-style-type: none"> • Airport, railway, and subway station early-warning threat detection • Customs and border patrol
CM Cross 科思创动	Employs deep-learning-powered image recognition to detect blood pressure, heart rate, and other physiological data.	<ul style="list-style-type: none"> • Customs and border patrol⁶⁷ • Early warning
EmoKit 翼开科技	EmoAsk AI Multimodal Smart Interrogation Auxiliary System detects facial expressions, body movements, vocal tone, and heart rate. ⁶⁸ Other products detect similar data for non-interrogation.	<ul style="list-style-type: none"> • Detecting and managing mental-health issues at medical institutions • Loan interviews at banks • Police-conducted interrogations⁶⁹
Joyware 中威电子	NuraLogix's DeepAffex is an image recognition engine that identifies facial blood flow (which is used to measure emotions) and detects heart rate, breathing rate, and 'psychological pressure'. ⁷¹	<ul style="list-style-type: none"> • Airport and railway station surveillance • Nursing • Psychological counselling
NuraLogix	Lower-level uses NuraLogix's polygraph tests. ⁷²	
Miaodong 秒懂	Relies on image recognition of vibrations and frequency of light on faces, which are used to detect facial blood flow and heart rate as a basis.	<ul style="list-style-type: none"> • Police interrogation
Sage Data 睿数科技	Public Safety Multimodal Emotional Interrogation System detects micro-expressions, bodily micro-expressions, and other physiological data.	<ul style="list-style-type: none"> • Police and court interrogations
Shenzhen Anshihua 深圳安世华	Emotion recognition product detects frequency and amplitude of light vibrations on faces and bodies.	<ul style="list-style-type: none"> • Early warning⁷⁶ • Prevention of crimes and safety
Taigusys Computing 太古计算	One product is referred to as a micro-expression-recognition system for Monitoring and Analysis of Imperceptible Emotions at Interrogation Sites, while others include 'smart prison' and 'dynamic emotion recognition' solutions. Taigusys claims to use image recognition that detects light vibrations on faces and bodies, as well as parallel computing. ⁷⁷	<ul style="list-style-type: none"> • Hospital use for detecting Alzheimer's, depression attacks⁷⁸ • Police interrogation of criminals⁷⁹ • Prison surveillance

As a society we need to agree the responsible limits on facial recognition and other high risk uses of AI technologies



Facial recognition technology: The need for public regulation and corporate responsibility

Jul 13, 2018 | [Brad Smith - President](#)



IBM quits facial-recognition market over police racial-profiling concerns

CEO writes to US Congress calling for 'national dialogue' about use in law enforcement

A group of humanoid robots are running on a track that features the United States and United Kingdom flags. The robot in the foreground is the largest and is running towards the viewer. It has a human-like face with a mustache. Behind it, several other robots of varying sizes are running in the same direction. The track is red with white and blue stripes, and the background is a bright, futuristic corridor with blue walls and white ceiling lights.


There is a leadership race for AI across nation states.

<https://www.gartner.com/technology/pressRoom.do?id=3872933>


Does this impact AI ethics?

China is an AI super-power with an ambition to be a global leader in AI by 2030

AI Super-Powers, Dr Kai Fu-Lee




“China is the Saudi Arabia in data”



“I left out Europe...”

“I left out Europe because I didn’t think there was a good chance for it to take even a so-called ‘bronze medal’ in this AI competition. European artificial intelligence is losing the race.”

Chinese structural advantages include data privacy regulations, public - private cooperation, and scale of data, capital, talent and market demand.

A portrait of Ursula von der Leyen, a German politician, is the background of the slide. She is shown from the chest up, wearing a light blue jacket, against a dark blue background with yellow stars.

“Create an AI ecosystem of trust.” Ursula von der Leyen, Davos Agenda Jan '21.

*“...Some of us are deeply **concerned** about the role which will be left for human beings in a **world run by AI.**”*

*“Others worry about the **serious effects that algorithms** can have on the health of our **democracies**. Who is taking the final decisions? Who is steering the flow of information?”*

“What we see through social media platforms seems real....we literally live in different worlds.”

*“Yes, algorithms can be a danger to our democracy. But they do not have to be...There must be at least **transparency** on how the algorithm works....*

*“For people to accept a role for AI in such decisions, they must be **comprehensible**.*

*And they must respect **people's legal rights**...we have to be able to examine the workings of the system and to ensure **human oversight**.*

*Our aim is to create an **AI ecosystem of trust**.*

**“Create an AI ecosystem of trust.” Ursula von der Leyen,
Davos Agenda Jan '21.**

“What sets Europe apart from competitors like China is not the size of our tech sector or how attractive we are for venture capital. What sets Europe apart from competitors is the fact that our values come first. Human beings come first.”

Thank you. Best Practice AI

Contact us:



Simon Greenman
Simon @ bestpractice . ai



@BestpracticeAi



+44 (0) 7 8 2 4 5 5 7 9 7 9
www.bestpractice.ai



Strategy, Business Planning & Education



AI Delivery
Support & Coaching



Organisational Design and
Capacity Building



AI Governance & Risk
Management





Investment Due Diligence


Further reading on AI governance

Ensure AI explainability, auditability within a governance framework

UK Information Commissioner's Office (ICO) Guidance



Model AI Governance Framework



(launched Davos 2020)