Using Big Data for Good (and not Evil)

AIEC

Duncan Ross
DataKind UK

@duncan3ross

DKUK

Not all data **is** big

Not all learning is deep

Not all decisions are fair

In this presentation…

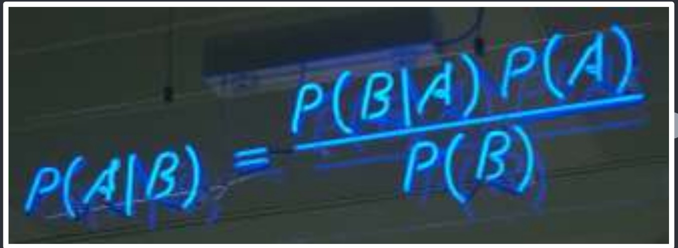

Limited formulae

No code

Some interaction…

DataKind is a global network

@duncan3ross

Recap: why do we use data science?

- People are inconsistent
- People have many biases
- Sometimes people are working in stressed situations
- People cannot process large volumes of data
- Decisions need to be made more rapidly
- Data science can provide better decisions
- Data science can provide new insights

What can you do?

- Make predictions
- Group things
- Gain insights
- Take decisions
- Do it at scale and speed

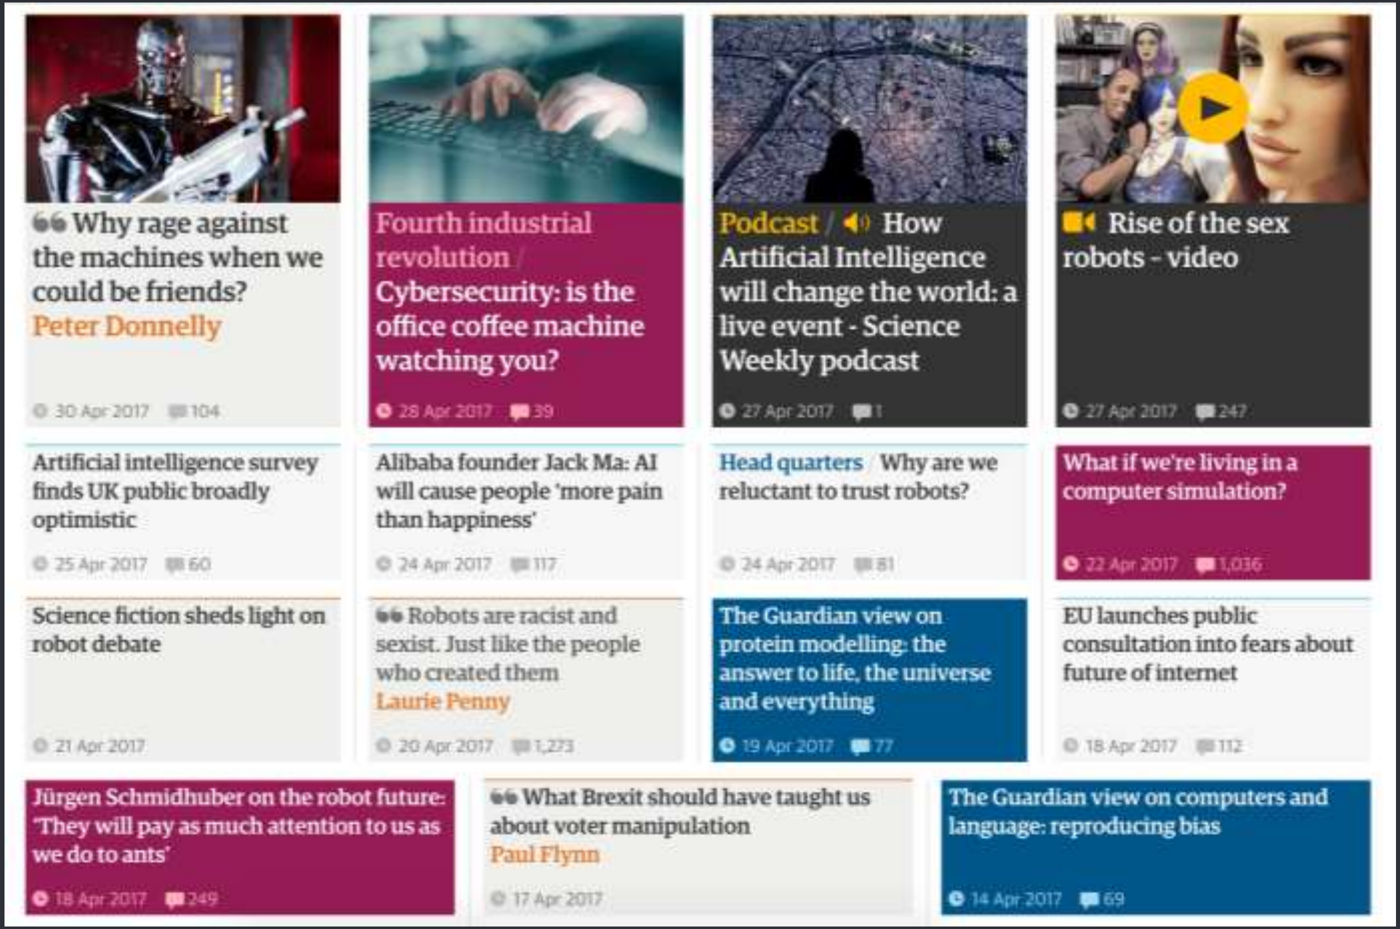
Why rage against the machines when we could be friends?
Peter Donnelly
30 Apr 2017 104

Fourth industrial revolution / Cybersecurity: is the office coffee machine watching you?
28 Apr 2017 39

Podcast / How Artificial Intelligence will change the world: a live event - Science Weekly podcast
27 Apr 2017 1

Rise of the sex robots - video
27 Apr 2017 247

Artificial intelligence survey finds UK public broadly optimistic
25 Apr 2017 60

Alibaba founder Jack Ma: AI will cause people 'more pain than happiness'
24 Apr 2017 117

Head quarters / Why are we reluctant to trust robots?
24 Apr 2017 81

What if we're living in a computer simulation?
22 Apr 2017 1,036

Science fiction sheds light on robot debate
21 Apr 2017

Robots are racist and sexist. Just like the people who created them
Laurie Penny
20 Apr 2017 1,273

The Guardian view on protein modelling: the answer to life, the universe and everything
19 Apr 2017 77

EU launches public consultation into fears about future of internet
18 Apr 2017 112

Jürgen Schmidhuber on the robot future: 'They will pay as much attention to us as we do to ants'
18 Apr 2017 249

What Brexit should have taught us about voter manipulation
Paul Flynn
17 Apr 2017

The Guardian view on computers and language: reproducing bias
14 Apr 2017 69

@duncan3ross

Stross-M-Banks Continuum

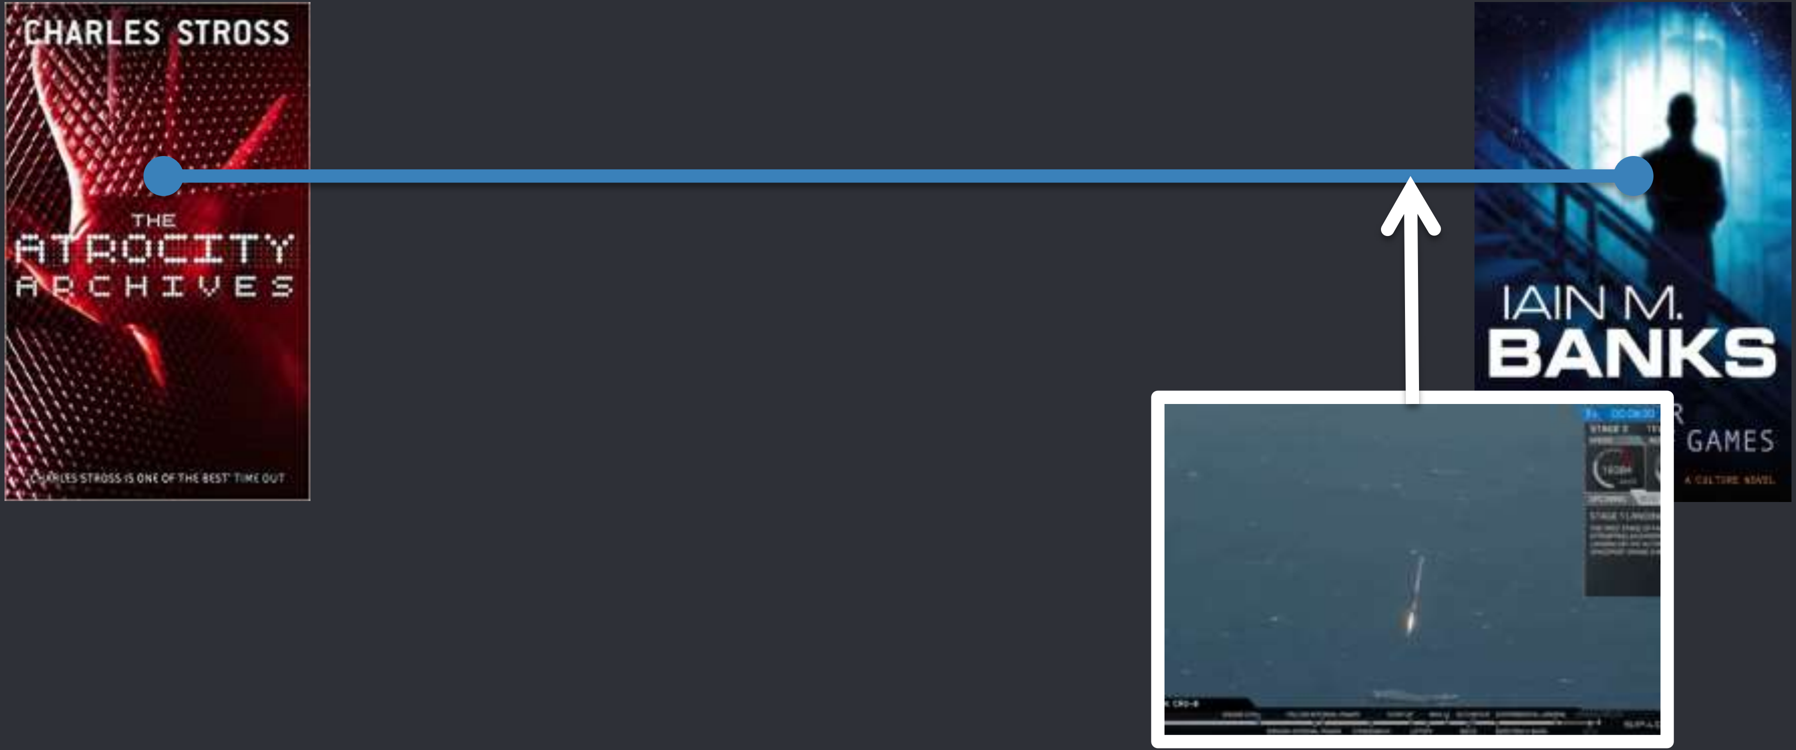

How is data being used in education?

- In a limited way
- Some signs of change
- What does that mean for diversity?

Invisible Aid

- USTC China
- Identify poor students through eating patterns
- > 60 meals < 200 Yuan
- Automatic 160 Yuan subsidy

Progress trajectory

- Academy Trust in UK
- Identify individual learning trajectories
- Aim to understand when to intervene

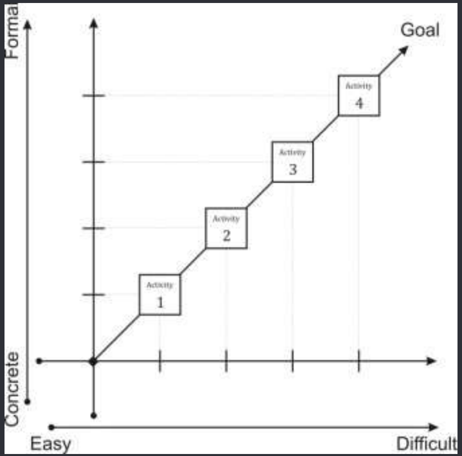

These both look good… so what’s the problem?

When algorithms go wrong…

- Racial bias in forecasting recidivism
- African Americans were more likely to be predicted as future criminals
- Impact on bail decisions

- Cathy O'Neill –
  Weapons of Math Destruction


BUSINESS INSIDER DE INTERNATIONAL

The racial bias that ProPublica found in a formula used by courts and parole boards to forecast future criminal behavior arises inevitably from the test's design, according to new research.

A white defendant.

The findings were described in scholarly papers published or circulated over the past several months. Taken together, they represent the most far-reaching critique to date of the fairness of algorithms that seek to provide an objective measure of the likelihood a defendant will commit further crimes.

Gefällt mir 0 Teilen 0

It’s not just the algorithm,
it’s what you do with it

When good algorithms go bad

Data science salary predictor:

“We created a basic, parsimonious linear model using the lasso with R^2 of 0.382. Most features were excluded from the model as insignificant”

70577 intercept
+1467 age
(per year above 18; e.g., 28 is +14,670)
-8026 gender=Female
+6536 industry=Software
(incl. security, cloud services)
-15196 industry=Education
-3468 company size: <500
+401 company size: 2500+
-15196 industry=Education
+32003 upper management
(director, VP, CxO)
+7427 PhD
+15608 California
+12089 Northeast US
-924 Canada
-20989 Latin America
-23292 Europe (except UK/I)
-25517 Asia

Source: O'Reilly 2015 Data Science Salary Survey
https://www.oreilly.com/ideas/2015-data-science-salary-survey

This algorithm is doing good!

- How much are data scientists paid?

–$23292
Europe

- I won't move there!
- I will vote to leave the EU!

70577 intercept
+1467 age
(per year above 18; e.g., 28 is +14,670)
-8026 gender=Female
+6536 industry=Software
(incl. security, cloud services)
-15196 industry=Education
-3468 company size: <500
+401 company size: 2500+
-15196 industry=Education
+32003 upper management
(director, VP, CxO)
+7427 PhD
+15608 California
+12089 Northeast US
-924 Canada
-20989 Latin America
-23292 Europe (except UK/I)
-25517 Asia

This algorithm is doing *more* good!

- How much are data scientists paid?

–$8026 gender=Female

- Better fix that inequality now we've noticed it!

70577 intercept
+1467 age
(per year above 18; e.g., 28 is +14,670)
-8026 gender=Female
+6536 industry=Software
(incl. security, cloud services)
-15196 industry=Education
-3468 company size: <500
+401 company size: 2500+
-15196 industry=Education
+32003 upper management
(director, VP, CxO)
+7427 PhD
+15608 California
+12089 Northeast US
-924 Canada
-20989 Latin America
-23292 Europe (except UK/I)
-25517 Asia

This algorithm is doing *evil*!

- How much are data scientists paid?

  –$8026 gender=Female

- I will use this to ensure that I don’t overpay women

70577 intercept
+1467 age
(per year above 18; e.g., 28 is +14,670)
-8026 gender=Female
+6536 industry=Software
(incl. security, cloud services)
-15196 industry=Education
-3468 company size: <500
+401 company size: 2500+
-15196 industry=Education
+32003 upper management
(director, VP, CxO)
+7427 PhD
+15608 California
+12089 Northeast US
-924 Canada
-20989 Latin America
-23292 Europe (except UK/I)
-25517 Asia

This algorithm is confusing

- How much are data scientists paid?

–$8026 gender=Female

- What if I'm giving advice to people who are hiring?

70577 intercept
+1467 age
(per year above 18; e.g., 28 is +14,670)
-8026 gender=Female
+6536 industry=Software
(incl. security, cloud services)
-15196 industry=Education
-3468 company size: <500
+401 company size: 2500+
-15196 industry=Education
+32003 upper management
(director, VP, CxO)
+7427 PhD
+15608 California
+12089 Northeast US
-924 Canada
-20989 Latin America
-23292 Europe (except UK/I)
-25517 Asia

Can we rely on Government and Laws?

Handbook of Robotics, 56th Edition, 2058 A.D.

- A robot may not injure a human being or, through inaction, allow a human being to come to harm.
- A robot must obey the orders given it by human beings except where such orders would conflict with the First Law.
- A robot must protect its own existence as long as such protection does not conflict with the First or Second Laws.

How did that work out?

Beyond the Law: Fairness

- Fairness isn't really a legal construct
- It certainly isn't a measure of a model's accuracy
- We call it colinearity, they call it intersectionality…

Challenges about fairness

- If you feed an algorithm racist data it will inevitably be racist
- Algorithmic unfairness is inherently worse than human unfairness
- Algorithms should be made transparent

DKUK

Challenges/claims

- If you feed an algorithm racist data it will inevitably be racist
- Possibly, but will it be more or less racist? Do we need perfectly unbiased algorithms?

- Algorithmic unfairness is inherently worse than human unfairness
- Nope.

- Algorithms should be made transparent
- Very difficult to achieve

A framework for using data: Baselines

- We need to understand the baseline
- We need to think about Type I and Type II errors
- We need to think about the cost of not doing something

DKUK

A framework for using data: Communication

- People are bad at understanding maths
- People are worse at understanding statistics
- Explicability vs accuracy
- Kahnemann: Thinking Fast and Slow

Some guidelines from MetroLabNetwork

- Engage
  - Internally
  - Externally (include skeptical partners!)

- Validate the model

- Open up about the model

https://metrolabnetwork.org/wp-content/uploads/2017/09/Ethical-Guidelines-for-Applying-Predictive-Tools-within-Human-Services_Sept-2017.pdf

Conclusions: what can we do to bridge the gap?

- Think about decisions
- Think before you analyse
- Understand how to communicate output
- Take personal responsibility

A data science pledge?

- *I will be **Aware** of the outcome and impact of my analysis*
- *I won't be **Arrogant** – and I will avoid hubris: I won't assume I should, just because I can*
- *I will be an **Agent for change**: use my analytical powers for positive good*
- *I will be **Awesome**: I will reach out to those who need me, and take their cause further than they could imagine*

Datakind.org.uk

@duncan3ross

@DataKindUK