Ethics and AI in medicine

Dr Richard Milne
Society and Ethics Research Group
Wellcome Genome Campus

Society and Ethics Research

Empirical social science research exploring the translation of genomics from bench to bedside and beyond

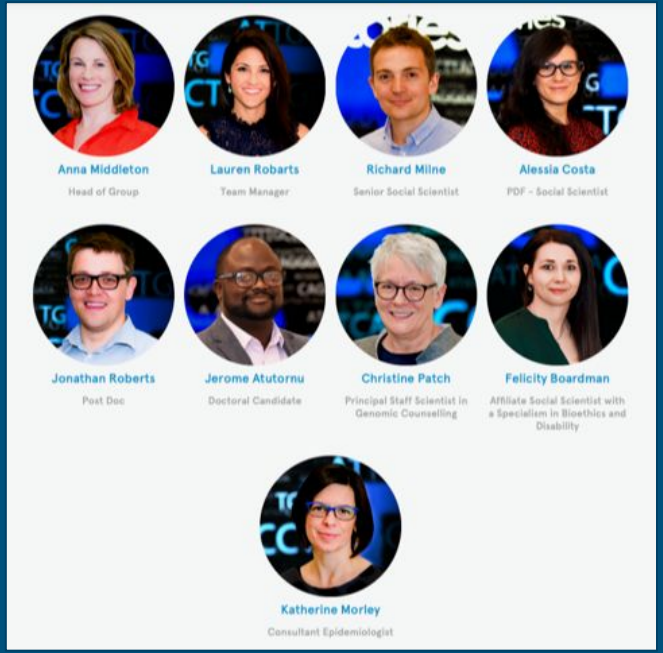

Disrupting the future of medicine?

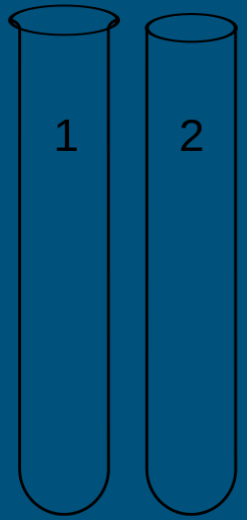

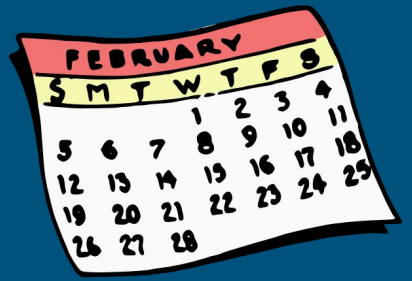

Wainberg et al., “Deep Learning in Biomedicine,” *Nature Biotechnology* https://doi.org/10.1038/nbt.4233.

The potential of AI

""the science of making machines do things that would require intelligence if done by people"

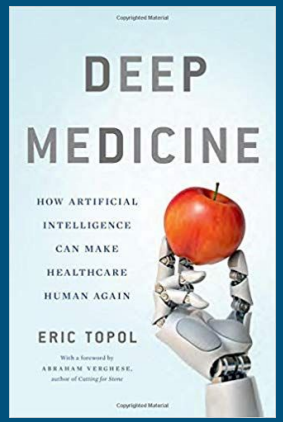

ESMO
Annals of Oncology 29: 1836–1842, 2018
doi:10.1093/annonc/mdy166
Published online 28 May 2018

ORIGINAL ARTICLE

Man against machine: diagnostic performance of a deep learning convolutional neural network for dermoscopic melanoma recognition in comparison to 58 dermatologists

H. A. Haenssle[1*,†], C. Fink[1†], R. Schneiderbauer[1], F. Toberer[1], T. Buhl[2], A. Blum[3], A. Kalloo[4], A. Ben Hadj Hassen[5], L. Thomas[6], A. Enk[1] & L. Uhlmann[7]

Diagnostics	Knowledge Generation	Public Health	System Efficiency	P4 Medicine
• Image Recognition e.g. • Symptoms Checkers and Decision Support • Risk Stratification	• Drug Discovery • Pattern Recognition • Greater knowledge of rare diseases • Greater understanding of casuality	• Digital epidemiology • National screening programmes	• Optimisation of care pathways • Prediction of Do Not Attends • Identification of staffing requirements	• Prediction of deterioration • Personalised treatments • Preventative advice

The challenges of AI

Societal and political concerns about data, algorithms and AI

How to maximise the benefits of the technology while responding to the ethical and social challenges it raises?

Semantics derived automatically from language corpora contain human-like biases

Aylin Caliskan,[1]* Joanna J. Bryson,[1,2]* Arvind Narayanan[1]*

"if we build an intelligent system that learns enough about the properties of language to be able to understand and produce it, in the process it will also acquire historical cultural associations, some of which can be objectionable."

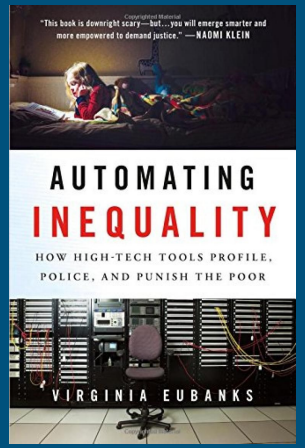

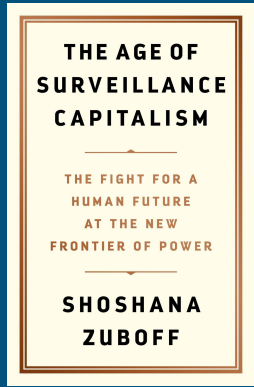

Machine Bias

There's software used across the country to predict future criminals. And it's biased against blacks.

by Julia Angwin, Jeff Larson, Surya Mattu and Lauren Kirchner, ProPublica
May 23, 2016

A commercial tool **COMPAS** automatically predicts some categories of future crime to assist in bail and sentencing decisions. It is used in courts in the US.

Prediction Fails Differently for Black Defendants

	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.)*

https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing

From Julia Stoyanovich GA4GH plenary 2019

‘Getting it right’

1. Respect for persons
2. Respect for human rights
3. Participation
4. Accounting for decisions

AI ethics guidelines

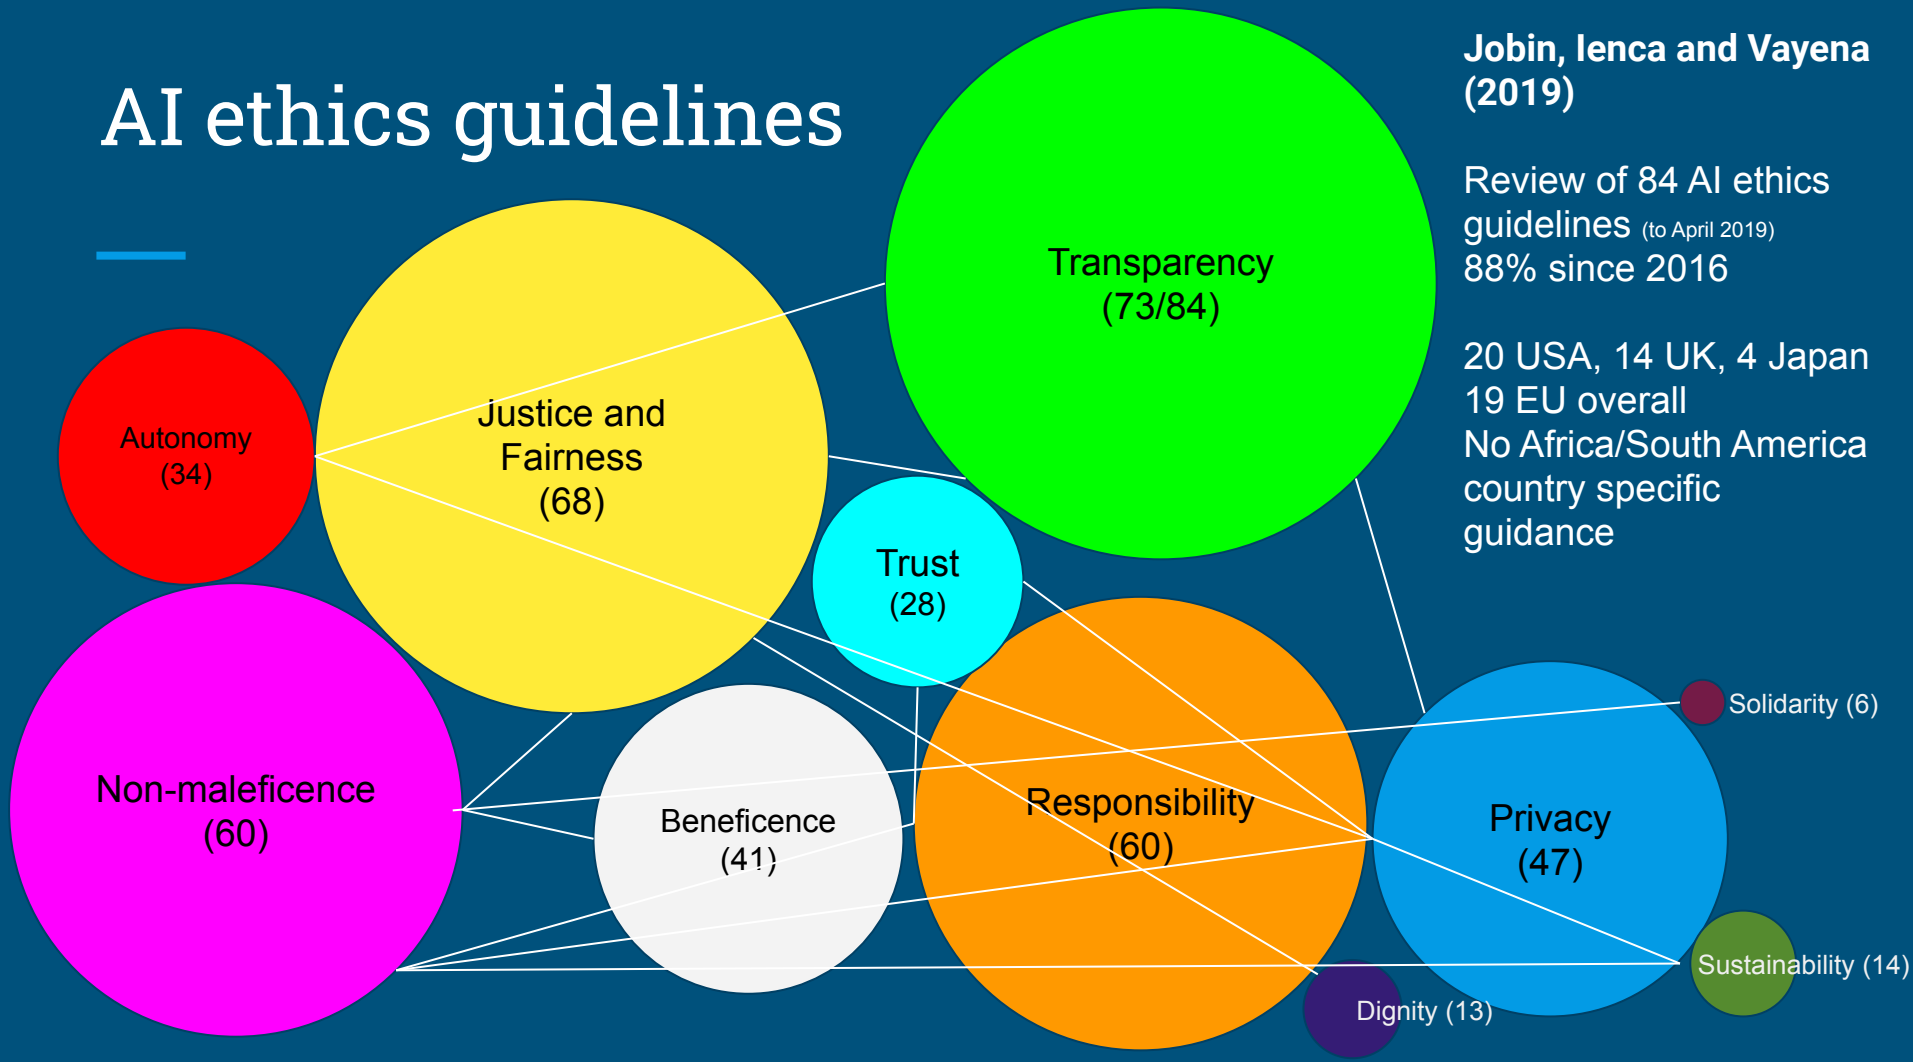

Jobin, Ienca and Vayena (2019)

Review of 84 AI ethics guidelines (to April 2019)
88% since 2016

20 USA, 14 UK, 4 Japan
19 EU overall
No Africa/South America country specific guidance

AI ethics guidelines

AI ethics guidelines

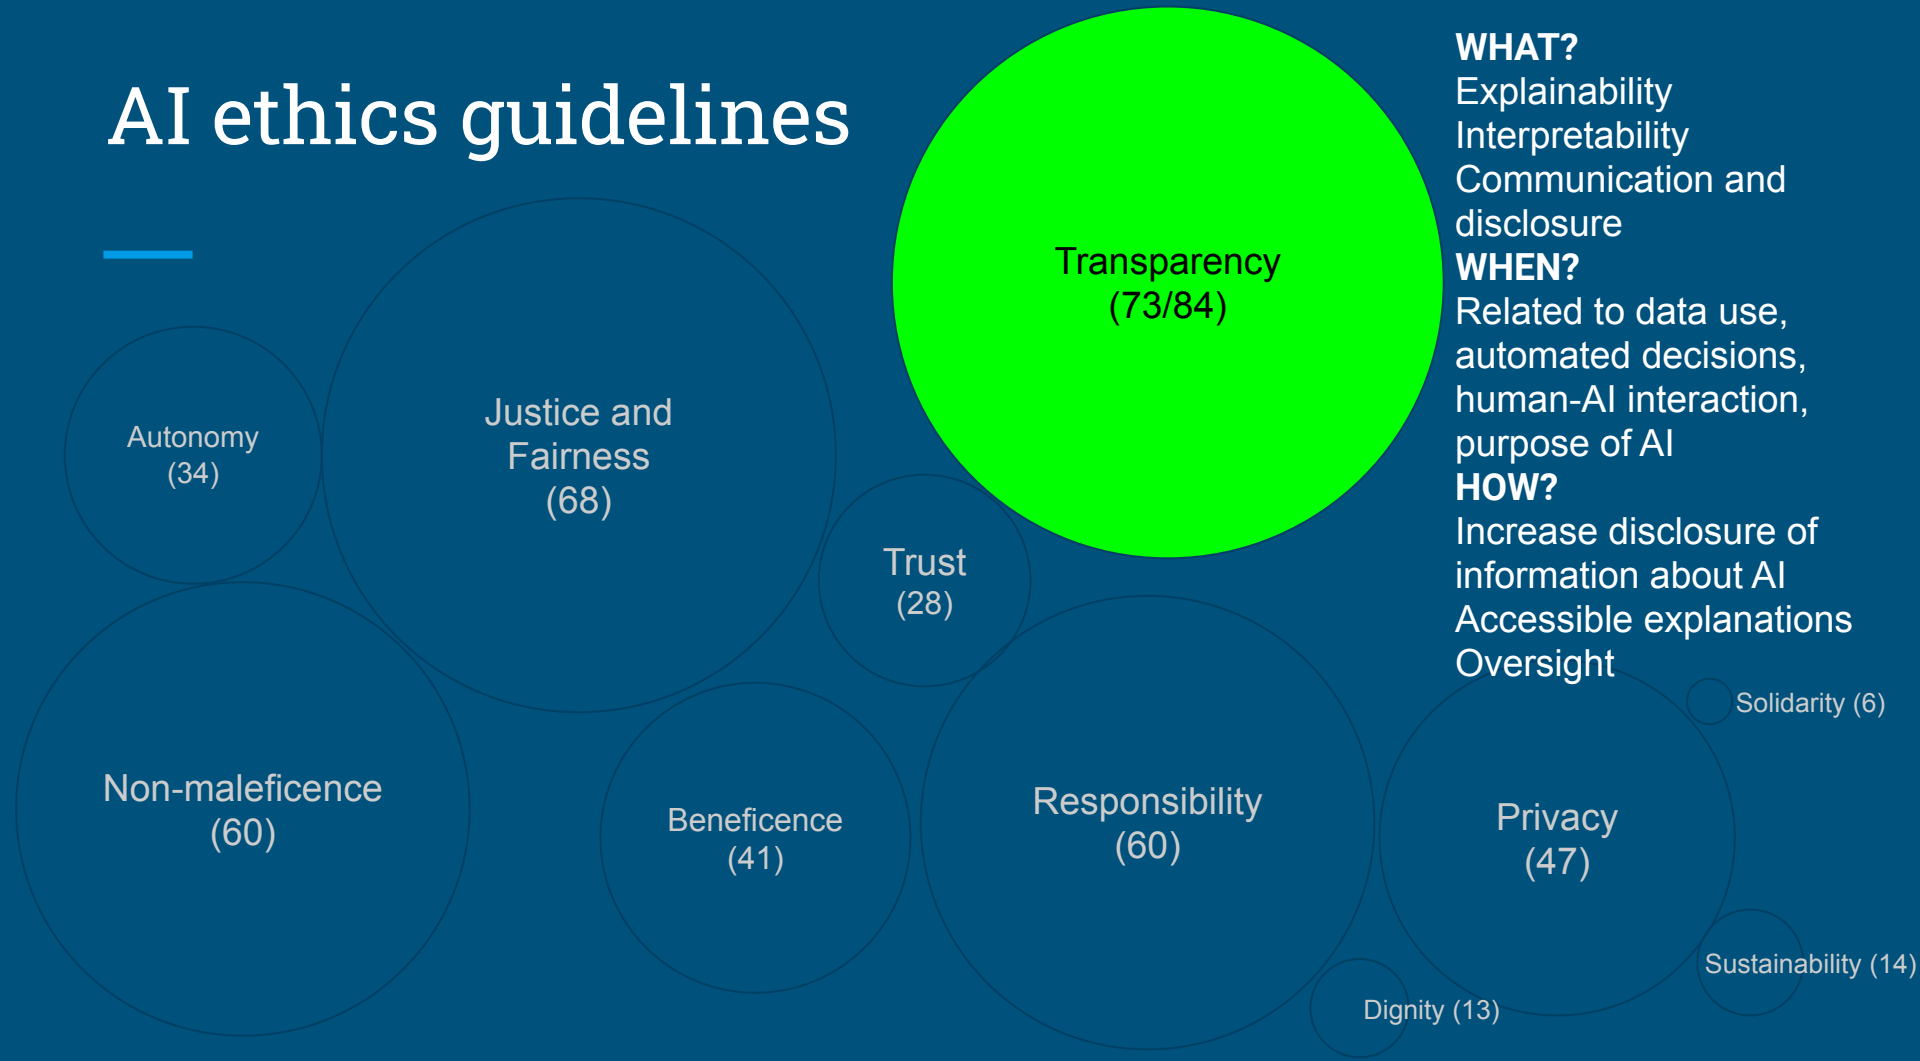

WHAT?
Explainability
Interpretability
Communication and disclosure

WHEN?
Related to data use, automated decisions, human-AI interaction, purpose of AI

HOW?
Increase disclosure of information about AI
Accessible explanations
Oversight

AI ethics guidelines

WHAT?
Prevention of unwanted bias and discrimination
Inclusion
Ability to challenge/remedy
WHEN?
Data acquisition and processing
HOW?
Technical solutions
Transparency
Monitoring
Law
Oversight

AI ethics guidelines

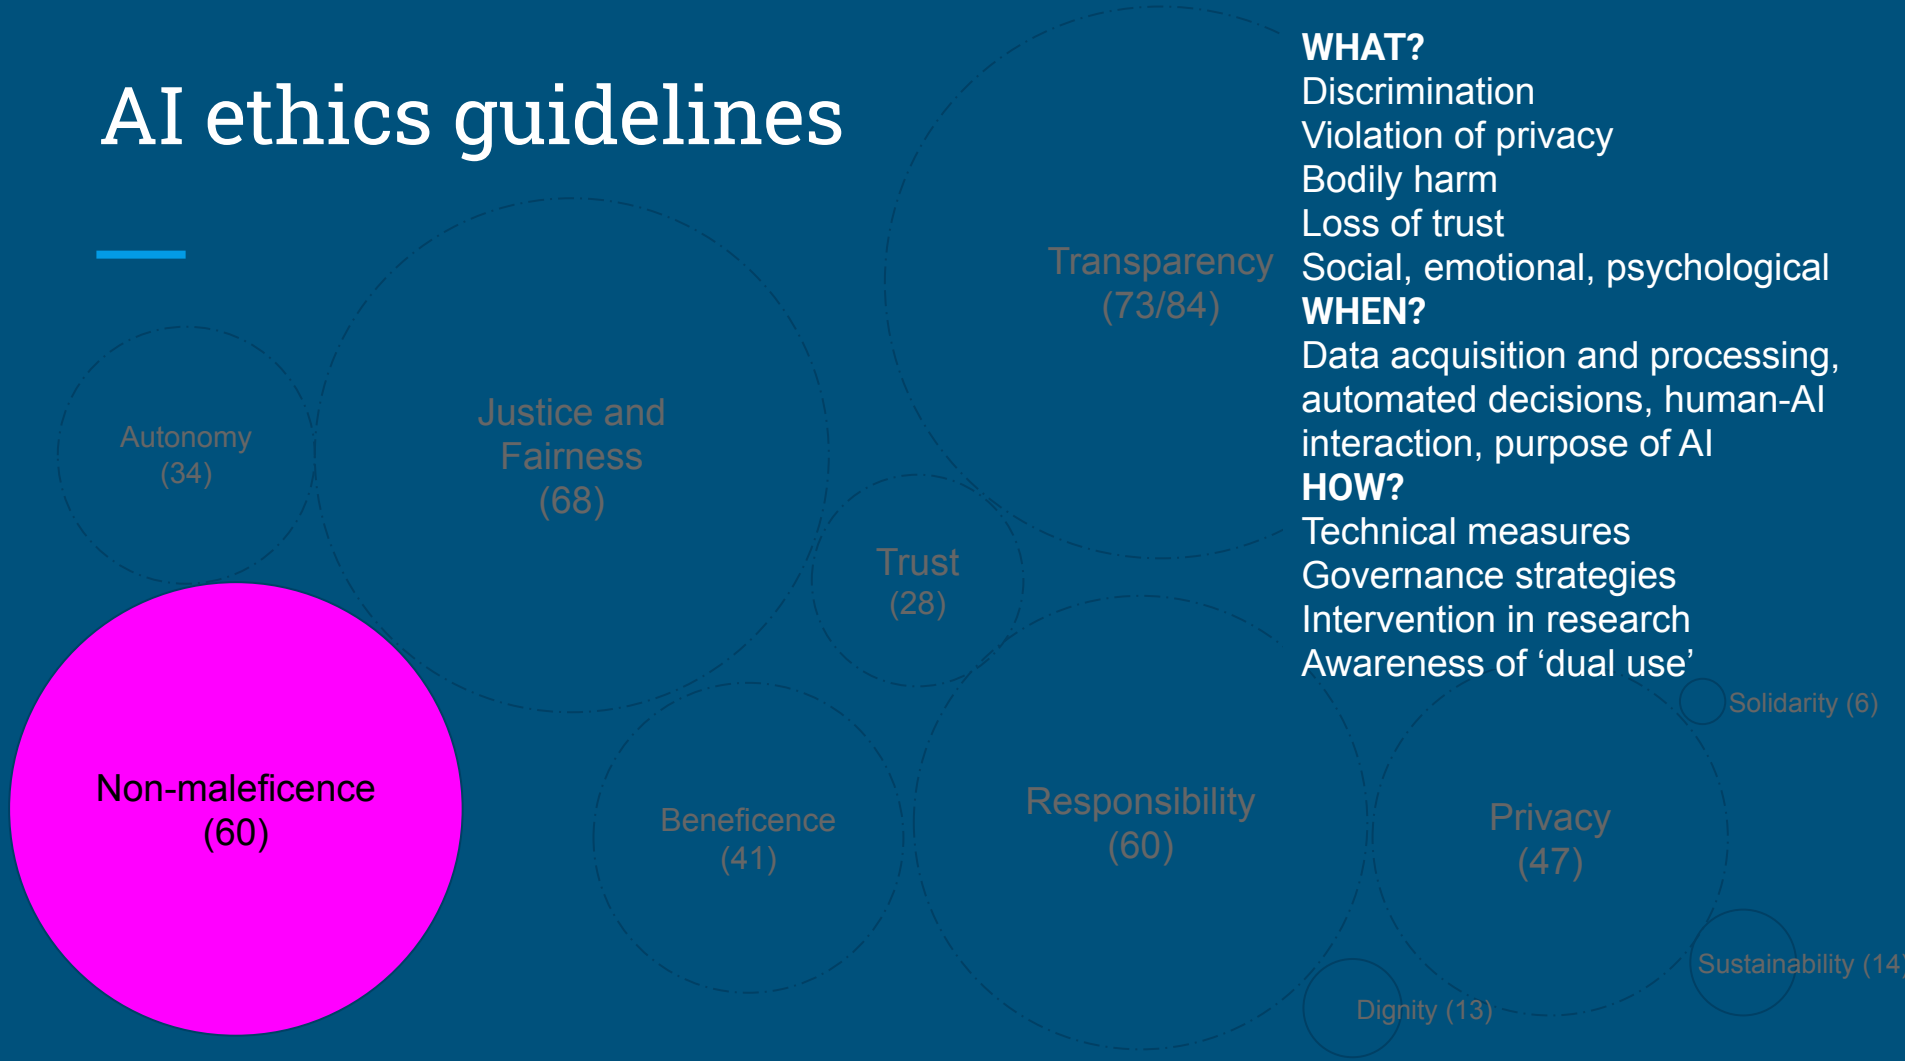

WHAT?
Discrimination
Violation of privacy
Bodily harm
Loss of trust
Social, emotional, psychological

WHEN?
Data acquisition and processing, automated decisions, human-AI interaction, purpose of AI

HOW?
Technical measures
Governance strategies
Intervention in research
Awareness of ‘dual use’

Case studies

Polygenic risk scores

Photographic phenotyping for rare diseases

Digital detection of cognitive decline

Polygenic risk scores

Potential for more accurate identification or stratification of individuals on the basis of risk for common conditions

- changing risk assessments and need to update (transparency)
- risk communication (non-maleficence)
- availability of datasets (justice)
- limited generalisability of PRS

Need more, better quality data on diverse populations and consideration of value and meaning of genetic information

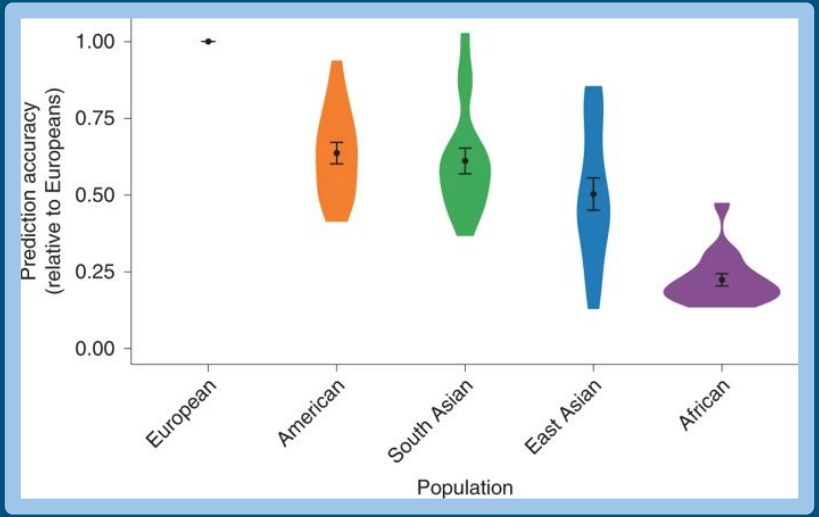

Prediction accuracy relative to European-ancestry individuals across 17 quantitative traits and 5 continental populations in the UKBB. (Martin et al. 2019)

Photographic phenotyping

Developments in sequencing facilitate accurate diagnosis of dysmorphologies

Rely on clinician expertise in identifying phenotypes

Machine learning algorithms can be used to detect facial features associated with intellectual disability and interpret VUS

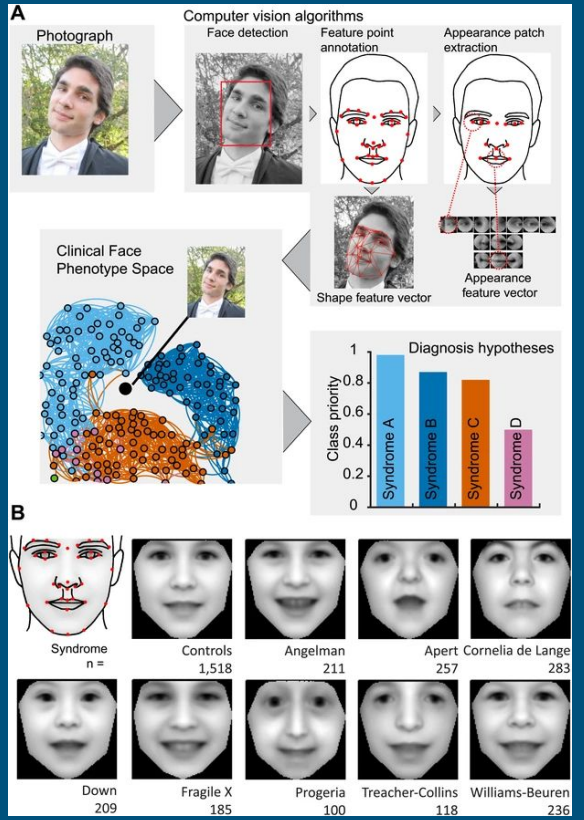

Ferry et al., eLife (2014) https://doi.org/10.7554/eLife.02020.

van der Donk et al.,*Genetics in Medicine* (2019) https://doi.org/10.1038/s41436-018-0404-y.

data-induced discrimination

- effective tool relies on representative dataset of photographs
- ethnic diversity in presentation
- potential for bias and uneven access

the management of incidental findings

- how to determine what IFs are potentially present in design?
- how to review and evaluate IFs?

commodification of phenotypic datasets

- data have value that developers may wish to protect
- phenotyping has public benefit - how to prevent data siloing

(Hallowell et al. 2019)

what is detected in whom, and what are the consequences?

- sensitivity of facial features
- potential for ‘off-label’ use
- legacy of stigma and discrimination based on photographic evidence
- reinforces social and racial stereotypes about face, race and intelligence

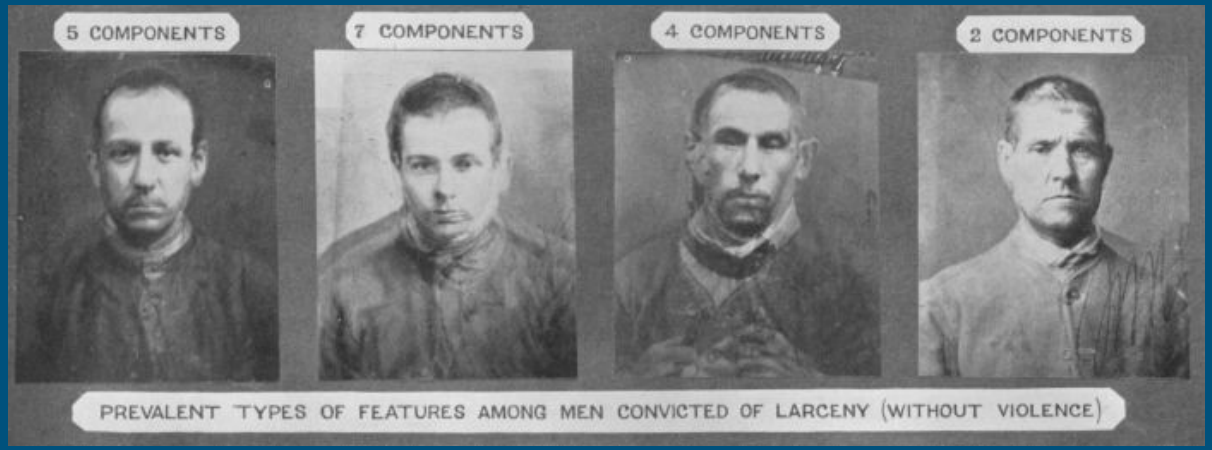

The dual use problem and the ‘double bottleneck’

“ethical mistakes or misunderstandings may lead to social rejection or distorted legislation and policies, which in turn may cripple the acceptance and advancement of data science” (Mittelstadt 2019)

Kate Crawford and Trevor Paglen, “Excavating AI: The Politics of Training Sets for Machine Learning (September 19, 2019) https://excavating.ai

Digital detection of cognitive decline

1. Many people with dementia do not have access to a ‘timely’ diagnosis

   Need to improve/extend ability to detect dementia

2. Repeated failure of Alzheimer’s drug trials attributed to ‘wrong people, wrong time’

   Identification of at-risk individuals and early detection of cognitive decline may enable targeted interventions

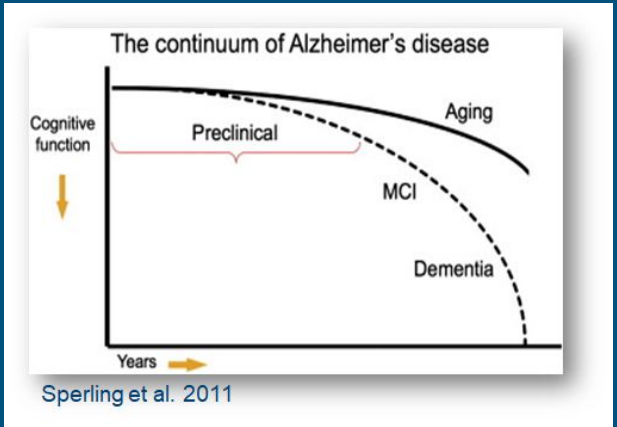

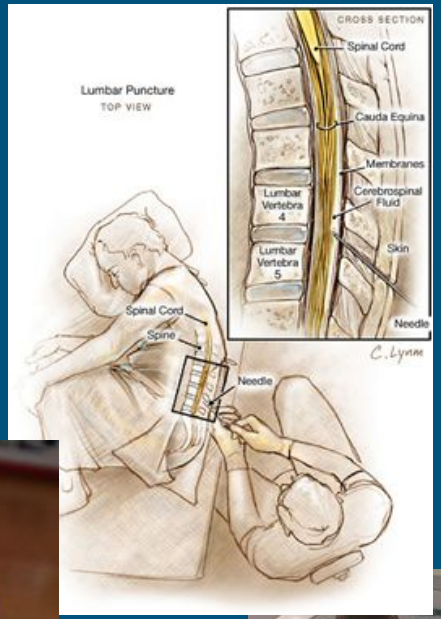
Lumbar Puncture
TOP VIEW
CROSS SECTION
Spinal Cord
Cauda Equina
Membranes
Cerebrospinal Fluid
Skin
Needle
Lumbar Vertebra 4
Lumbar Vertebra 5
Spinal Cord
Spine
Needle
C. Lynn

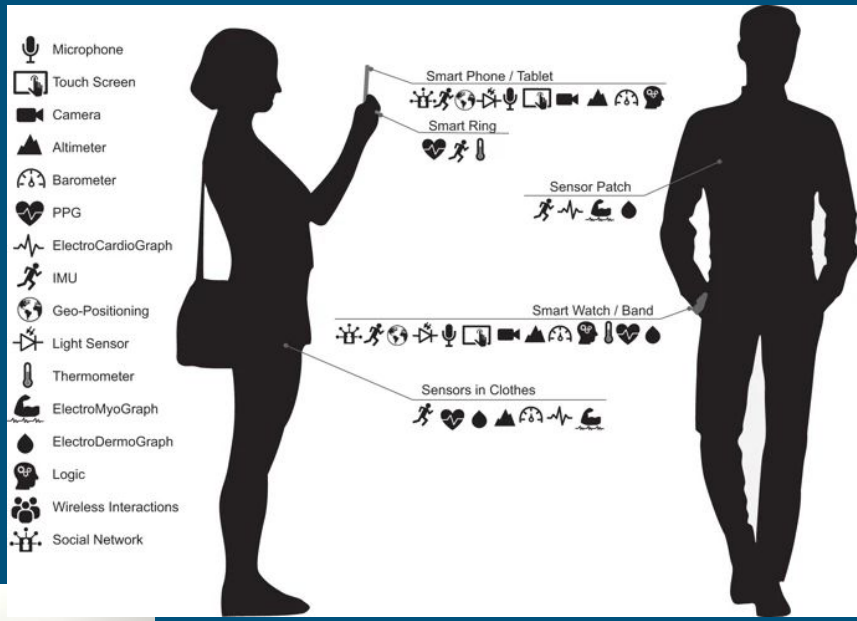
Microphone
Touch Screen
Camera
Altimeter
Barometer
PPG
ElectroCardioGraph
IMU
Geo-Positioning
Light Sensor
Thermometer
ElectroMyoGraph
ElectroDermoGraph
Logic
Wireless Interactions
Social Network
Smart Phone / Tablet
Smart Ring
Sensor Patch
Smart Watch / Band
Sensors in Clothes


FILL TO

New data, new analytics

Implementation of ML in image recognition, analysis of large-scale behavioural data, new forms of data collection

Active/passive measures

Improve translation from cell/animal/lab to human/clinic

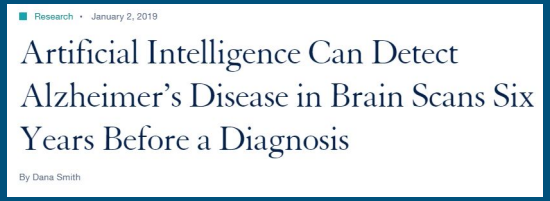

Cognetivity Neurosciences

FACE PRINT™

SPACE (Stakeholder Perspectives on social and ethical Aspects of digital Cognitive Evaluation)

What ethical challenges are emerging in practice?

How does something get identified as an ‘ethical’ question?

What gets done about it?

What do members of the public think about these uses of data about them?

Data ethics

“A lot of people are data hoarding - it's not always immediately obvious until you ask for data, but a lot of people have said know when we've asked the data to do machine learning on some some people have said no because they're publishing on the data themselves” (clinician researcher)

“It [ethics] really has all been around data. So kind of who owns the data, where's the data being stored who's processing it and some concerns about that” (academic researcher)

Transparency

How to be transparent about changing uses of data?

I think they play out differently because because for a company … We're constantly developing on a product which means we're capturing new types of data every day. That already complicates things for us. The data is also super important for feeding back into our research and development process of the product. So there are there more users of the data for for business than there is for an academic.

Responsibility

Who does detection?

How do we justify decisions?

SLIDES PRESENTING UNPUBLISHED DATA - PLEASE DO NOT CITE WITHOUT PERMISSION

“I think the problem is what we're technically capable of doing is not something that we necessarily can do ... I mean, I can produce algorithms which could diagnose dementia possibly as accurately as a clinician

… If we're thinking practically then obviously you get into a hotter and hotter water the more responsibility you take away from the clinician, even if it leads to better and better patient outcomes.” (clinician researcher)

Justice

“we know that the diagnosis varies on the basis of ethnicity and things like that and and some people believe that there are gender biases and ethnic biases in the [human] diagnosis and if that's true and we develop a 100% accurate classifier, I'm concerned that we're essentially, at extreme you could say we've developed a racist algorithm for example” (clinician researcher)

Timing

“We should probably also be targeting people before they go to before they make a decision about whether to go to the doctor or not.“ (clinician researcher)

Timing

And regulation

Once you're talking about people who are healthy at least by clinical standards who don't have any diagnosis, then you can is like that it's almost like a consumer business right? I mean you can you cannot you cannot claim anything clinical, but maybe you're not trying to do anything clinical. You're not trying to treat a diagnosis, you're trying to help them reduce the risk of even getting a diagnosis. (company researcher)

Privacy

Who detects detection?

"A system and method that enables a person to unobtrusively assess their cognitive function from mobile device usage." (Mindstrong)

"The invention embodiments provide the potential for an automated, seamless and non-intrusive detection of different diseases." (Fujitsu)

“the system may monitor household occupants' movement patterns and compare these movement patterns with those associated with Alzheimer's disease … the system may infer a higher probability that the household occupant has the disease … The inference may be reported and/or recorded for subsequent use in the system.”

(US patent US9872088B2)

Google Home

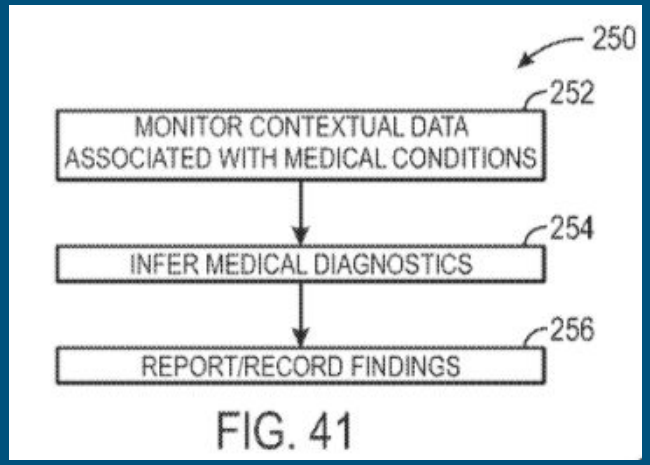

FIG. 41

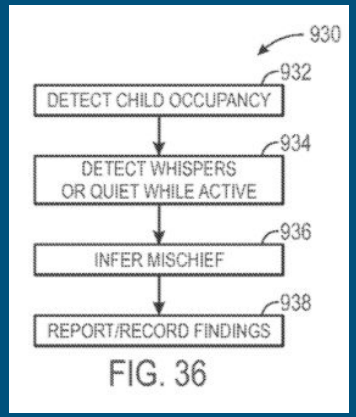

FIG. 36

Ethics and early detection of cognitive change

- ‘Data ethics’
  - Access, ownership, sharing etc.
- ‘Detection ethics’
  - Risk communication, fear
- ‘Algorithm ethics’
  - Transparency
  - Justice/Fairness
  - Responsibility
  - Privacy
- Trust and commercialisation

Conclusions

Applications of ML in medicine raise a range of questions related to data, detection and algorithms

We are interested in talking to people working with ML/deep learning etc. about what they're doing, what challenges they encounter and how to incorporate societal and ethical considerations at early stages of research

- Help us understand the field and the challenges it raises
- Hopefully help researchers anticipate (and avoid) emerging ethical concerns

Disrupting the future of medicine?

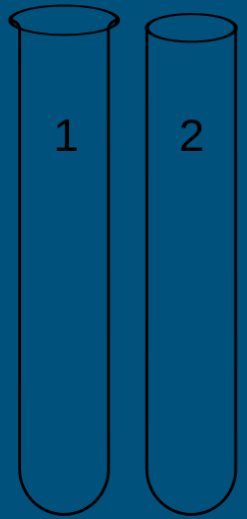

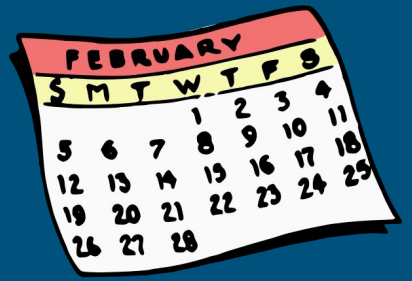

deep genomics

Wainberg et al., “Deep Learning in Biomedicine,” *Nature Biotechnology* https://doi.org/10.1038/nbt.4233.

SOCIETY AND ETHICS RESEARCH
SEMINAR SERIES

SOCIETY+
ETHICS
RESEARCH

From Genomes to GPS: socially responsible research and the future of data-driven medicine

Dr Richard Milne | Society and Ethics Research

Monday 9 December 12:30-13:30
Kendrew lecture theatre, EBI South building

CONNECTING
SCIENCE

References

Morley et al. *Artificial Intelligence: How to get it right* NHSX https://www.nhsx.nhs.uk/assets/NHSX_AI_report.pdf

Anna Jobin, Marcello Ienca, and Effy Vayena, 'The Global Landscape of AI Ethics Guidelines', Nature Machine Intelligence 1, no. 9 (September 2019): 389–99, https://doi.org/10.1038/s42256-019-0088-2.

Camillia Kong, 'Ethical Dangers of Facial Phenotyping through Photography in Psychiatric Genomics Studies', Journal of Medical Ethics 45, no. 11 (1 November 2019): 730–35, https://doi.org/10.1136/medethics-2019-105478.

Brent Mittelstadt, 'The Ethics of Biomedical "Big Data" Analytics', Philosophy & Technology 32, no. 1 (1 March 2019): 17–21, https://doi.org/10.1007/s13347-019-00344-z.

Nuffield Council on Bioethics, 'The Collection, Linking and Use of Data in Biomedical Research and Health Care: Ethical Issues' (London: Nuffield Council on Bioethics, 2015).

Michael Wainberg et al., "Deep Learning in Biomedicine," Nature Biotechnology 36, no. 9 (October 2018): 829–38, https://doi.org/10.1038/nbt.4233;

Aylin Caliskan, Joanna J. Bryson, and Arvind Narayanan, "Semantics Derived Automatically from Language Corpora Contain Human-like Biases," Science 356, no. 6334 (April 14, 2017): 183–86, https://doi.org/10.1126/science.aal4230;

Roos van der Donk et al., "Next-Generation Phenotyping Using Computer Vision Algorithms in Rare Genomic Neurodevelopmental Disorders," Genetics in Medicine 21, no. 8 (2019): 1719–25, https://doi.org/10.1038/s41436-018-0404-y;

Quentin Ferry et al., "Diagnostically Relevant Facial Gestalt Information from Ordinary Photos," eLife, June 24, 2014, https://doi.org/10.7554/eLife.02020;

Nina Hallowell, Michael Parker, and Christoffer Nellåker, "Big Data Phenotyping in Rare Diseases: Some Ethical Issues," Genetics in Medicine 21, no. 2 (February 2019): 272–74, https://doi.org/10.1038/s41436-018-0067-8.