

Ethics in Data Science

Dr Ben Collier



Overview

- What are research ethics and why do we have them?
- Basic practical aspects of doing principled, ethical digital research
- The new wave – data science after the AI turn
- Wide ethical issues of digital research

Why do we have research ethics?

- Ethics is a long-standing concern of scientific practice
- Historically a concern of professional communities
- Some active regulation of the harms of industrial progress (pollution, Factory Acts)
- Middle and second half of the 20th Century see a huge public reaction against the harms of science and technology

Science historically bound to society, and hence to harm

- How we produce knowledge
- Involvement in and justification of colonial regimes
- Eugenics
- Edinburgh's own history of involvement in colonialism, slavery, eugenics
- Weapons and defence research
- Science inextricable from wider social and political relations, forms of society



Tuskegee Syphilis Study

- 1932 Alabama, USA
- US Public Health Service
- 600 African American men selected – told being treated
- Actually – goal is to observe progress of untreated syphilis
- Active attempts made to ensure men not treated, including telling doctors not to treat them
- Continues up to 1972
- National Research Act 1974 – informed consent
- Last survivor died in 2009



Nazi science and the Nuremburg Code

- Eugenic and industrial, chemical science deeply bound to the Nazi project and implicated in the Holocaust
- Nuremberg War Crimes Tribunal
- 10 standards for research on human subjects
 - Informed consent
 - Wider social benefit and necessity
 - Prior knowledge justifies performance
 - Avoid unnecessary suffering and injury
 - No expectation of death or disabling injury
 - Risk versus reward
 - Protections against remote risks of injury
 - Qualified persons
 - Opt-out and cessation
 - Physician prepared to cease in good faith
- Directly shapes the Belmont Report – respect, beneficence, justice



Ethics and social research

- Anthropology and colonial legacies
- Eugenics and criminology – unethical forms of knowledge

Harms of 20th Century sociology

- Harms from reporting
- Vulnerable participants
- Genuinely informed consent
- Outcomes for groups
- Ethnography/intimacy

Developments in sociological ethics

- Feminist and civil rights movements
 - Beyond informed consent – active participants
 - Shaping and ownership of voice
 - Framing of questions
 - Sharing of benefits
- Critique of colonial and geopolitical dynamics of ‘Western’ research, and role of scientists and publics from the Global South
- Global and local political economies of research – who benefits from, sets the boundaries and questions for, carries out research? Who is the ‘subject’ and who is the ‘object’?

Quantitative

- Population level surveys and administrative data
- Full consent given or administrative consent
- Purposes and re-use (usually – scientific research, not marketing)
- Most data highly identifying
- Can still be extremely sensitive

Qualitative

- Interviews and observation
- Consent usually negotiated directly
- Power dynamics (esp in professional or institutional settings - can they disagree?)
- Discomfort and direct harm
- Quote review
- Right to withdraw
- Analysis?
- We are not journalists – we have different (often more restrictive) duties to participants

Practical guide – university research ethics

- Identify harms and risks
 - To you
 - To them
 - To the University
 - To the profession
 - To wider society
- Propose mitigations
- Positive rather than negative
- Wider goals as well as specific harms
- Compliance and process as well as ethics and moral values

Anonymity

- Statistical de-anonymization esp for small numbers
- Sparsity
- Table privacy
- Subtraction attacks
- Small counts
- Data combinations
- Differential privacy
- Location data

	£20k	£30k	>£40k
A	2	10	0
B	20	23	0
C	15	30	0
D	0	0	1
Fail			

Some methods for protecting datasets

- Remove direct identifiers
- Band data to categorical
- Paraphrase or generalize qual data
- Recoding categorical data
- Top and bottom coding
- Perturbation and record swapping
- Thresholds for display of small counts

Benefit sharing and extractivism

- Who actually benefits from the research?
- Do you have funding? Can you compensate directly?
- Are you asking them to do work?
- Relative power dynamics
- How might benefits be shared?
- Can be – communicating findings
- Participant ‘voice’
- Policy impacts

Examples

- Statistical studies – developing models of association
- Qualitative interviews with computer hackers
 - Online
 - Age?
 - Withdrawal?
 - Privacy?
 - Vulnerability?

Data science and ethics

– the basics

- Who cares what those sociologists are up to? I work with DATA!
- Large-scale collections of the traces of human life
- Similarities to ethnographic collections
- New types of data – location, health, large amounts of text, activity, behaviour
- Very difficult to fully anonymise at scale
- Large amounts of data in the public realm and available for private purchase
- OSI skills – deduction, subtraction, boiling down anonymity sets
- Huge datasets, often sparse, easily linked – trivially deanonymisable



Privacy and computing

- If you build systems or research pipelines – security
- Ethical and legal duty to manage data responsibly and securely
- Data Management Plan
- Encryption and secure storage
- Data chain of custody – where does it stay? Who has access? How long before it's deleted?
- Intermediate outputs, code, weights?

Is Reddit the same as a high street?

- Street photography (and how is that different from research)?
- Do/can people meaningfully consent?
- Counting numbers of people versus detailed covert photography
- What sort of street is it?
- Do users take meaningful precautions?
- Financial/labour activities – production of content, rights?
- Large-scale scraping decreasingly used in real-name environments



Data, collection and consent

- Assumed for many years in the data science boom that data on public platforms was more-or-less public
- Use of a platform implies consent
- But – context? How do people actually use the platform?
- Clandestine collection

Terms of service – (when) do we care?

- Robots.txt
- Set of legal(ish) statements about proper use of the platform
- What you agree to as a user
- Major platforms now usually direct you to the API
- But – holding power to account?

Facebook is obstructing our work on disinformation. Other researchers could be next

Laura Edelson and Damon McCoy

The company's hostility to academic scrutiny limits our ability to understand how the platform amplifies political falsehoods



📷 'If Facebook has its way, there will be no independent research of its platform.' Photograph: Jaap Arriens/NurPhoto/Rex/Shutterstock

Last week, Facebook disabled our personal accounts, obstructing the research we lead at New York University to study the spread of disinformation on the company's platform. [The move](#) has already compromised our work - forcing us to suspend our investigations

Analytics and ethics

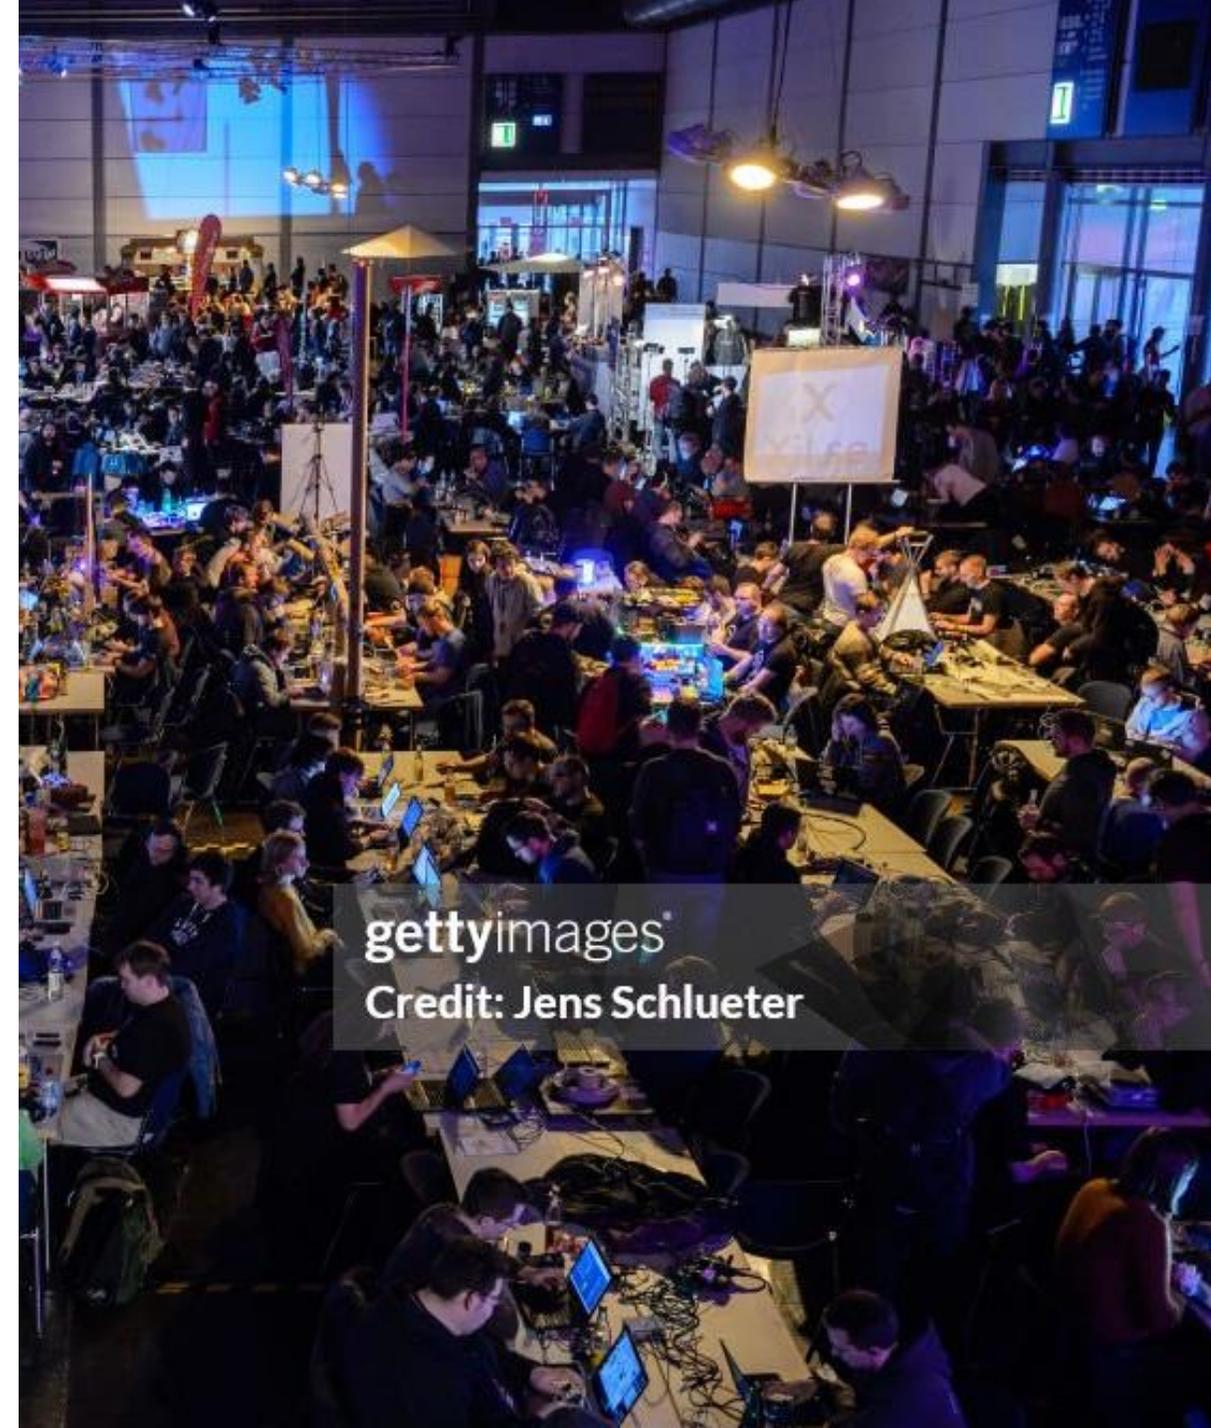
- Can we find out things about people that they don't know about themselves?
- If (for example) you are building a research algorithm that tries to detect early onset illness from people's twitter posts, do you have a duty to tell them?

Practical approaches – mixed methods

- Netnography
- What do people seem to expect? How many people actually see a given post?
- Use norms

Example: CrimeBB

- Cambridge Cybercrime Centre
- More than 100 million posts scraped from Clear and Darknet cybercrime communities, Discord and Telegram channels
- Clandestine collection
- Legally unenforceable ToS
- Community members anonymous (but lots of personal information)
- Idea is to share data widely with academics under strict legal agreements



Experiments

- This has so far been about collecting data – observing social life
- But what about actively changing it?
- Example – Facebook experiment
- Alternatives – lab-based, simulation

Facebook reveals news feed experiment to control emotions

Protests over secret study involving 689,000 users in which friends' postings were moved to influence moods

[Poll: Facebook's secret mood experiment: have you lost trust in the social network?](#)



📷 Activists and politicians called Facebook's experiment 'scandalous', 'spooky' and 'disturbing'.
Photograph: Dado Ruvic/Reuters

GDPR and other major EU/UK legislation

- Several legal developments in the past ten years
 - Regulated in the UK by the Information Commissioner's Office
 - Rights as data subjects
 - Personal data – relates to living identifiable people
 - Data processing must be lawful, fair, and transparent
 - Basis:
 - Research - 'task in the public interest'
 - Commercial – 'legitimate interest'
 - Consent
 - Data can be stored for research indefinitely where legitimate justification
 - Special category data – must be necessary for research in the public good
- **Must:**
 - Not cause substantial damage/distress
 - Not make decisions that affect individuals on basis of research data
 - Respect data minimization
 - Anonymise/pseudonymize where possible
 - Protect privacy and security
 - Establish public interest through peer review

GenAI

- Initial collections arguably illegal and unethical
- Opposite of minimization by design
- Sustainability
- Large-scale changes to platforms and data science
- Platforms increasingly full of generated content
- New possibilities for experimental and quasi-experimental designs
- Ethical questions about representation and rights of data subjects – ‘epistemic justice’
- Ethics of releasing possibly-harmful new products to a wide consumer market
- But how do you know what the harms are going to be?
- Mitigations – guardrails, testing, non-use, minimisation

Changing professional norms in computational sociology and data science

- Increasing reluctance to mass scrape data (and legal barriers)
- Not just analysis and data – significant potential harms from new technological affordances (and new legal duties on you e.g. under OSA)
- Potential step-changes in scale
- But – ethical concerns with pure simulation, consequences for participant voice, democratic role of research

Ethics beyond the experiment

Sustainability

- Social, environmental, and financial costs of compute
- Duty to minimize?
- Efficient, lean analytical approaches?
- Use the simplest tool that works well



Working with and for big tech

A moment of reckoning has arrived for the West.

Our culture has fallen into shallow consumerism while abandoning national purpose. Too few in Silicon Valley have asked what ought to be built — and why.

We did.

We built Palantir to ensure America's future, not to tinker at the margins. On the factory floor, in the operating room, across the battlefield — we build to dominate.

Join us.


ALEXANDER C. KARP
CEO & CO-FOUNDER
PALANTIR TECHNOLOGIES INC.


PALANTIR.COM
PALANTIR TECHNOLOGIES
INC.



Bloomberg

Zuckerberg Says Most Companies Need More 'Masculine Energy'



Using the tools of big tech

- Important questions about who is shaping the results your tools give you
- Politics, weights, bias -> how can you evaluate your own pipelines? Unit testing, performance testing?
- Security of the data tools you use – where is your data going?
- Engage in depth with your microdata and what each stage of your workflow is producing
- Wider ethics – commitment to e.g. Open Source Software



Working for (or against) the police and the state

- Ultimate aims of research
- Predictive policing – bias, risk, harm
- Defence research
- Controls over future uses of your research?
- Resistance research – hacks, leaks, data science (Forensic Architecture, DDoSecrets)



What is research ethics?

- A set of personal reflective practices
- Part of being a member of a profession
- Compliance with the processes and procedures of an organization
- Much bigger set of practices and values mediating the conditions in which knowledge about the world is produced and what that knowledge is then used to do
- NOT – a single hurdle at the start of a project