

What are research ethics?

Daniel R. Thomas[1] **Sara Correia**[2] Helena Webb [3]

[1]Cambridge Cybercrime Centre, Department of Computer Science and Technology, University of Cambridge, UK

[2]University of Swansea, UK

[3]Department of Computer Science, University of Oxford, UK

SecHuman 2018

Daniel: 5017 A1EC 0B29 08E3 CF64 7CCD 5514 35D5 D749 33D9

In this session we are going to think about what research ethics are. We will look at the principles and issues and use a range of case studies to apply those ideas. In the next session we will look more at the process by which we ensure that research is ethical.

Format

1. Warm up (10 minutes)
2. Short lecture (20 minutes).
3. Group discussions (20 minutes)
4. Plenary discussion (30 minutes)
5. Feedback and consolidation of key issues (10 minutes)

Warm up

In groups of up to 5 people discuss for 5 minutes:

- Where have you seen ethical issues arise in research?
- What kinds of things go wrong?
- What are research ethics for?

We will then feed back for 5 minutes.

Think before doing research that might effect people

- Consider the ethics carefully
- Describe your considerations in any publications
- “humans” not “human participants” or “human subjects”
- Seek review as necessary (see next session)

Describing your ethical considerations in publications means we can learn from each other, and reviewers can properly assess your research.
Think about how humans (and the environment more generally) will be affected by your research, not just about the humans who might be direct participants or subjects of your research.

We had questions about ethics in our research[1]

- UDP honeypot DDoS sensors
- Developed statistical method for estimating coverage
- Verified using leaked booter databases...



[1] Daniel R. Thomas, Richard Clayton, and Alastair R. Beresford. 2017. 1000 days of UDP amplification DDoS attacks. In *APWG Symposium on Electronic Crime Research (eCrime)*. IEEE, Scottsdale, AZ, USA.

Why did we start down this road? Well we had questions about ethics in our research, which I presented here last year.
Some of these leaked booter databases contain the names and email addresses of children and an association with illegal activity. We did not analyse that part of the data but as part of sorting the data out so that I could process it my eyes saw it, and I have endevoured not to remember it.
We have the leaked databases of over 35 booters and we can share at least some of those with other researchers.

Ethics are norms of conduct

- Distinguish between acceptable and unacceptable behaviour
- Guidance in Menlo Report[2]
- Enforced by REBs and Program Committees (e.g. IMC[3])
- Minimise harm, maximise benefit.

[2] David Dittrich, Michael Bailey, and Erin Kenneally. 2013. Applying ethical principles to information and communication technology research: A companion to the Menlo Report. Tech. rep. U.S. Department of Homeland Security, (Oct. 2013).

[3] Mark Allman and Vern Paxson. 2007. Issues and etiquette concerning use of shared measurement data. In *Proceedings of the 7th ACM SIGCOMM conference on Internet measurement*. ACM, 135–140.

What are ethics? In this work we consider ethics to be norms of conduct.

Menlo report

- Respect for persons
- Beneficence
- Justice
- Respect for law and public interest

Individuals should be treated as autonomous agents and persons with diminished autonomy should be given additional protection.
Minimise possible harms, and maximise possible benefits. The researcher should also use safeguards against potential harms.
Risks and benefits should be distributed fairly and not on the basis of protected characteristics such as race, or other characteristics that correlated with protected ones.
In general ethical research conforms to applicable laws in relevant jurisdictions. Research should always be in the public interest. Addition- ally, research should be open, transparent, reproducible and peer-reviewed.

Observation or intervention[4]

These sessions focus on ethics of observational studies. There are many more considerations for intervention based studies.
For careful reasoning on the ethics of intervention based research see:

https://civilservant.io/

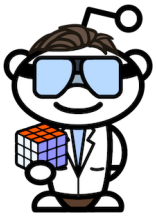

https://www.youtube.com/watch?v=RLzRtDp7ES0

[4] J. Nathan Matias and Merry Mou. 2018. CivilServant: Community-led experiments in platform governance. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing System*. ACM Press, 1–13. ISBN: 978-1-4503-5620-6.

Our experience is with observational studies, there is no clear boundary between observational and intervention based studies but as the research tends towards being more intervention based greater care is required due to the greater risk of direct harm.
The CivilServant work is a good example of carefully designed intervention work looking at questions like how moderation tools for online communities should work.

We should consider these ethical issues[5]

- Identification of stakeholders
- Informed consent (where possible)
- Identify harms
- Safeguards
- Justice
- Public interest

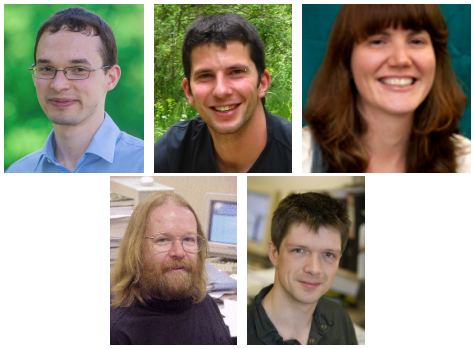

[5] Daniel R. Thomas, Sergio Pastrana, Alice Hutchings, Richard Clayton, and Alastair R. Beresford. 2017. Ethical issues in research using datasets of illicit origin. In *Proceedings of the Internet Measurement Conference (IMC)*. ACM. London, UK, (Nov. 2017).

From our analysis of existing guidance, the issues which should be considered when proposing to use data of illicit origin are:
The stakeholders should be identified to support the analysis of the potential harms and benefits of the research.
If you don't realise that law enforcement might be investigating the system you are studying then you might not take precautions to avoid interfering with their investigation.
We will talk more about informed consent in the next session. Informed consent means that the participants understand the research and the potential harms and benefits well enough to make an informed choice.
The potential harms should be identified. Could this censorship measurement study cause people to be arrested because they appeared to be accessing censored pages?
Having identified potential harms, what can we do about them? Do we anonymise the results in publications to protect individuals?
The research does not unfairly advantage or disadvantage any group. If research only benefits rich people who can afford to make use of the results, is that good enough?
The research is not just interesting to you but is also of public interest.
Definitions taken from our paper which forms the basis of much of this session.

Legal issues may arise

- Computer misuse
- Copyright
- Data privacy (e.g. GDPR)
- Terrorism
- Indecent images
- National security
- Contracts

IANAL. Just pointers: not even the beginning of advice.

Misuse or abuse of computers such as unauthorised access and the use of malware or 'dual use' tools.
Can you share the data you are using for your research if you lack the copyright/database rights or if there are trade secrets?
Is personal data involved? Are IP addresses personal in your context (IPs of servers are probably not but IPs of web browsers probably are). If personal data is involved what is your grounds for processing? Public interest research?
In some jurisdictions (e.g. UK) it may be an offence to fail to report terrorist activity. If you are deliberately doing research into terrorism then check your institution's guidance to make sure you won't get arrested.
Possession of indecent images of children is an offence i many jurisdictions. Where it is possible that your research might accidentally pick up indecent images you need procedures in place to detect that and respond appropriately without any human having to look at any images.
If you do research using leaked classified materials, will your institution delete your research?
If you promised not to do something and then you do it then that is problematic.

AT&T iPad users database breach

- "Researchers" from Goatse Security found vulnerability in AT&T website
- Obtained email addresses for 114 000 iPad users
- Shared vulnerability, exploit, and email addresses with third parties
- Did not report vulnerability to AT&T
- Went to jail

Unethical and illegal.

An example of what can go wrong if computer misuse issues (and ethics in general) are not properly considered. Lack of justice in the work as the method was designed to advantage the researchers notoriety at the expense of everyone else.

Illegal but ethical research?

- *Mens rea*: lack of criminal intent
- Not in the public interest to prosecute (perhaps ask the prosecutor first)
- REB approval! REB protection?
- Get the law fixed

In some cases if the researcher can demonstrate lack of criminal intent then a criminal prosecution cannot succeed; REB approval may be a useful way to demonstrate this.
This is an especially generic defence, but it is uncertain one. The last defence between you and court. Several researchers have conducted research after first obtaining approval from the prosecutor so that they knew they would not be prosecuted. However, you may need to obtain approval from prosecutors in many jurisdictions depending on the work in question.
If you get REB approval then perhaps the University lawyers will b defending you rather than attacking you.
If you are doing research which is ethical but illegal then you should talk to your elected representative about fixing the law.

Bad and good justifications

Bad:

- Not the first
- Public data

Good:

- No additional harm
- Fight malicious use
- Necessary data

Norms change, just because someone did research like this before does not mean you can do it now.
Researchers may develop or apply new techniques to public data that, for example, deanonymise these data, and this may cause harm. Clearly not every use of public data is good and so we need to consider whether our use of the data is good.
While harm might have occurred at some point, the analysis we do does not cause any additional harm and does have a public benefit.
Bad people are already using this data to to bad things ad we need to use this data to counter them.
There is no way of doing this research without using this data and this research has a clear public benefit.

Patreon database leak[6]

- Researching crowdfunding on Patreon by scraping the website (complete scrape?)
- Patreon is compromised and full database + source code leaked
- Authors decided there would be additional harm
  - Private vs. public data
  - Legitimize criminal activity
  - Violate privacy

[6] Nathaniel Poor and Roei Davidson. 2016. The ethics of using hacked data: Patreon's data hack and academic data standards. Data and Society. Tech. rep. Council for big data, ethics and society, (Mar. 2016), 1–7. Retrieved Sept. 28, 2017 from http://bdes.datasociety.net/council-output/case-study-the-ethics-of-using-hacked-data-patreons-data-hack-and-academic-data-standards/.

The authors decided against using the leaked data for their research even though it would have been useful, but they did get to write a paper about why they didn't use it.
Papers explaining why some work was not done are very useful and we should write more of them. Similarly for "We did this research but the PC rejected it as unethical" papers.

Safeguards against potential harm

- Secure storage
- Privacy
- Controlled sharing

Encrypt data at rest, servers in a secure location, strong access controls, up-to-date software.
Identities of individuals not revealed in publications.
Shared with other academic researchers under suitable contracts/agreements but not publicly.
WIP paper on this topic.

Potential harms

- Illicit measurement
- Potential Abuse
- DeAnonymisation
- Sensitive Information
- Behavioural Change
- Researcher Harm

Research involves illicit activities.
Research results could be used by malicious actors to cause additional harm, for example by means of designing evasive malware or updating password cracking policies.
Research on these data can be used to de-anonymise or re-identify people or networks. Also, identification of group of individuals may raise ethical concerns such as discrimination or violence to- wards identified groups.
These data contains sensitive and private information, which can be used to harm natural persons. For example, if the user password from one service is leaked, their credentials to other services can be compromised due to password reuse.
The research can change the behaviour of the stakeholders of these data, which may have negative consequences. For example, a market vendor can provide fake information if she knows that she is being measured.
The research can lead to the re- searchers being prosecuted by law enforcement, since these data may include illegal material. Researchers could be threatened by criminals, e.g. in underground forums [72], or by state or industry actors that dislike the work. There may also be a risk of emotional trauma to researchers if they come across distressing content, such as pornography or violence, during the work.

Researcher Harm

- Researchers are participants
- Ethics also protects researchers
- Maltese journalist Daphne Caruana Galizia killed by car bomb used Panama Papers data also used by researchers[7]
- Kenichiro Okamoto cybersecurity expert was killed by an Internet troll[8]

[7] Joseph Borg. 2017. Daphne Caruana Galizia obituary. *The Guardian*, (Nov. 21, 2017). Retrieved July 10, 2018 from http://www.theguardian.com/media/2017/nov/21/daphne-caruana-galizia-obituary.

[8] Jake Adelstein. 2018. A cyber-security expert trolled the trolls. Then one killed him. *The Daily Beast. world*, (June 29, 2018). Retrieved June 29, 2018 from https://www.thedailybeast.com/a-cyber-security-expert-trolled-the-trolls-then-one-killed-him.

Benefits provided

- Reproducibility
- Uniqueness
- Defence Mechanisms
- Anthropology & Transparency

The research can be reproduced and compared with other approaches.
The data used is unique (not otherwise obtainable) or historical (can no longer be obtained by other means).
Work allows the development of defence mechanisms such as anti-malware tools or better password policies.
Provides the ground truth on the real behaviour of human beings which other methods could only reveal in a filtered or biased way. For example, the real behaviour of password creation. Provides transparency into the behaviour of governments or corporations and so provides checks and balances on their power.

Safeguards, harms, and benefits[9]

Safeguards:

Privacy: No individuals identified

Controlled Sharing: Under legal agreement

Potential harms:

Sensitive Information: Names, email addresses, criminal activity

Behavioural Change: Don't advertise (working) booters

Benefits:

Uniqueness: Ground truth on booter activity

Anthropology & Transparency: Real behaviour of booters

[9] Daniel R. Thomas, Richard Clayton, and Alastair R. Beresford. 2017. 1000 days of UDP amplification DDoS attacks. In *APWG Symposium on Electronic Crime Research (eCrime)*. IEEE, Scottsdale, AZ, USA.

Lets do ethical research

- Think about the ethics at an early stage
- Take appropriate safeguards
- Write about it in publications

Group discussions

- Groups of up to 5 people
- Make notes on your case study:
  1. key ethical issues
  2. the harms that resulted
  3. could harms have been prevented?

Plenary discussion

Feedback from each group on the case studies they looked at.

Thank you! Questions?

Daniel R. Thomas Daniel.Thomas@cl.cam.ac.uk
https://www.cl.cam.ac.uk/~drt24/
@DanielRThomas24
5017 A1EC 0B29 08E3 CF64 7CCD 5514 35D5 D749 33D9

Sara Correia S.Correia@swansea.ac.uk
https://www.cyberlawandsociety.org/
@SGCorreia

Helena Webb Helena.Webb@cs.ox.ac.uk
https://www.cs.ox.ac.uk/people/helena.webb/
@EthicsWildfire

Daniel Thomas is supported by the EPSRC [grant number

References I

[1] Jake Adelstein. 2018. A cyber-security expert trolled the trolls. Then one killed him. *The Daily Beast. world*, (June 29, 2018). Retrieved June 29, 2018 from https://www.thedailybeast.com/a-cyber-security-expert-trolled-the-trolls-then-one-killed-him.

[2] Mark Allman and Vern Paxson. 2007. Issues and etiquette concerning use of shared measurement data. In *Proceedings of the 7th ACM SIGCOMM conference on Internet measurement*. ACM, 135–140.

[3] Joseph Borg. 2017. Daphne Caruana Galizia obituary. *The Guardian*, (Nov. 21, 2017). Retrieved July 10, 2018 from http://www.theguardian.com/media/2017/nov/21/daphne-caruana-galizia-obituary.

References II

[4] David Dittrich, Michael Bailey, and Erin Kenneally. 2013. Applying ethical principles to information and communication technology research: A companion to the Menlo Report. Tech. rep. U.S. Department of Homeland Security, (Oct. 2013).

[5] David Dittrich and Erin Kenneally. 2012. The Menlo Report: Ethical principles guiding information and communication technology research. Tech. rep. U.S. Department of Homeland Security, (Aug. 2012).

[6] J. Nathan Matias and Merry Mou. 2018. CivilServant: Community-led experiments in platform governance. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing System*. ACM Press, 1–13. ISBN: 978-1-4503-5620-6.

References III

[7] Nathaniel Poor and Roei Davidson. 2016. The ethics of using hacked data: Patreon's data hack and academic data standards. Data and Society. Tech. rep. Council for big data, ethics and society, (Mar. 2016), 1–7. Retrieved Sept. 28, 2017 from http://bdes.datasociety.net/council-output/case-study-the-ethics-of-using-hacked-data-patreons-data-hack-and-academic-data-standards/.

[8] Daniel R. Thomas, Richard Clayton, and Alastair R. Beresford. 2017. 1000 days of UDP amplification DDoS attacks. In *APWG Symposium on Electronic Crime Research (eCrime)*. IEEE, Scottsdale, AZ, USA.

References IV

[9] Daniel R. Thomas, Sergio Pastrana, Alice Hutchings, Richard Clayton, and Alastair R. Beresford. 2017. Ethical issues in research using datasets of illicit origin. In *Proceedings of the Internet Measurement Conference (IMC)*. ACM. London, UK, (Nov. 2017).