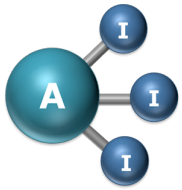

Artificial Intelligence and Augmented Intelligence for Automated Investigations for Scientific Discovery

The Physical Sciences Data-Science Service (PSDS)

Patterns

Failed it to Nailed it: Responsible Data Management: Legal & Ethical Aspects
19/11/2020
Online Event

Dr Samantha Kanza & Dr Nicola Knight
University of Southampton

12/03/2021

AI3SD-Event-Series:Report-21

Failed it to Nailed it: Responsible Data Management: Legal & Ethical Aspects
AI3SD-Event-Series:Report-21
12/03/2021
DOI: 10.5258/SOTON/P0034
Published by University of Southampton

Network: Artificial Intelligence and Augmented Intelligence for Automated Investigations for Scientific Discovery
This Network+ is EPSRC Funded under Grant No: EP/S000356/1

Principal Investigator: *Professor Jeremy Frey*
Co-Investigator: *Professor Mahesan Niranjan*
Network+ Coordinator: *Dr Samantha Kanza*

An EPSRC National Research Facility to facilitate Data Science in the Physical Sciences: The Physical Sciences Data science Service (PSDS)
This Facility is EPSRC Funded under Grant No: EP/S020357/1

Principal Investigator: *Professor Simon Coles*
Co-Investigators: *Dr Brian Matthews, Dr Juan Bicarregui & Professor Jeremy Frey*

Contents

1 Event Details

Title	Failed it to Nailed it: Responsible Data Management: Legal & Ethical Aspects
Organisers	AI^3 Science Discovery Network+, Patterns Journal & Physical Sciences Data-Science Service
Dates	19/11/2020
Programme	AI3SD Event Programme
No. Participants	41
Location	Online Event
Organisation / Local Chairs	Dr Samantha Kanza & Dr Nicola Knight

2 Event Summary and Format

This event was the third of the 'Failed it to Nailed it' online data seminar series. The event was hosted online through a zoom conference. The event ran for approximately 3 hours in an afternoon session.

There were three talks given on different aspects of responsible data management including the ethical dimensions of research data management and different types of legislation with respect to personal and non personal data. These talks covered some of the key concepts that are central to responsible data management alongside some practical methods of carrying these out. The talks were followed by an interactive breakout discussion on the moral and ethical considerations surrounding different aspects of the data lifecycle. The discussions were facilitated by using a set of Moral IT cards[1] developed by researchers at the University of Nottingham.

3 Event Background

This event is part of the 'Failed it to Nailed it' data seminar series. This event series, currently comprised of four online events is a collaboration between AI^3 Science Discovery Network+, Patterns & the Physical Sciences Data-Science Service (PSDS). This event series follows on from a data sharing survey that was undertaken earlier in 2020. Each event in the series handles a different aspect of dealing with data aiming to educate and inform researchers about how to work well with their data, as well as encouraging discussion along the way. Following on from these events the organisers hope to be able to organise more face-to-face events in 2021 which will expand this event series.

Ethical and legal aspects of data sharing are fundamental elements that need to be considered in any data-driven research, irrespective of domain or the type of data. In these events we want to encourage researchers to consider these as essential aspects of planning their initial research and data activities, rather than something that is considered post project or tacked on later as an afterthought. This event aimed to provide an introduction to different elements of ethics and legislation that pertain to data sharing, with suggestions for best practices and the opportunity to discuss these different aspects with peers from a range of domains.

[1] https://lachlansresearch.com/the-moral-it-legal-it-decks/

4 Talks

4.1 Ethical data management - balancing individual privacy and public benefit - Zosia Beckles (University of Bristol)

https://orcid.org/0000-0003-4352-2266

Figure 1: Zosia Beckles

The full video of Zosia's talk can be viewed here: https://youtu.be/J9kWkzK83i4 [1]

Zosia Beckles is an experienced information professional with a background in health informatics and research data management in the higher education and research sectors. She currently works with Library Research Support at the University of Bristol, providing training and support to academics in research metrics/bibliometrics and sensitive research data management, including the development of a new dataset disclosure risk assessment service to enable safe publication of sensitive data.

This talk from Zosia focused on ethical data management and how we can balance the individual right to privacy against the public benefit of sharing data, which is a key concept across many areas of research. The topics covered in Zosia's talk included: the big picture - what are the risks and rewards of sharing data, what can happen if things go wrong, regulatory frameworks, different ways to manage the risk and some best practice guidelines and tools.

Why do we want to share data?

Zosia begins by talking about the benefits that we can get from sharing our data:

- Maximise the value of research data: it takes a lot of effort and time to collect, both on the behalf of researchers but also research participants.
- Reduce the burden on participants: reusing data can reduce the burden.
- Improve the scientific record: making sure that we are keeping track of not only the methods and results that come out of studies but also the data itself that underpins those results.
- Verification and replication: data sharing can help combat the replication crisis by allowing other researchers to verify findings.

What are the risks?

The main risk area highlighted by Zosia is participant privacy. By releasing a data set containing information there is the risk of identification of a participant, this could occur through a

variety of different methods.

Identification from within a dataset

- Direct identification: inclusion of direct identifiers in the dataset e.g. names, addresses, email addresses
- Indirect identification: combinations of indirect identifiers e.g. age, gender, general geographic location, employment history etc., which could be enough to identify an individual when combined

Data Linkage outside of the dataset

- Dataset Linkage: creating links between entirely separate published datasets. An example of this [2] is the release of Netflix 'anonymous' subscriber data (times / dates for watching content + an anonymised identifier) being combined with public IMDB reviews and ratings for TV and movies, which allowed the identification of viewers through linking the time/date data even though the Netflix data was 'anonymous'.
- Information Linkage: even when there is not another dataset available data may be picked up from reports or papers that supplement the information given in the dataset. An example given is a dataset that doesn't include geographic location being combined with a methodology that describes where the participants were recruited from. This information could increase the risk of a participant's identity being revealed.

What happens when things go wrong?

There are a wide variety of things that could potentially go wrong through poor data management. Zosia focused first on the worst case scenario, this is a data breach, where information is disclosed that leads to the identification of specific participants within a research study. This may lead to harm to the individual participant but also to significant fines for the parties responsible for the breach of information. In a wider sense this can also damage the reputation of the researchers and have a negative impact on future research studies as participants may be less willing to be involved.

Other areas were also shared by Zosia where there may be issues if researchers have not properly thought about their data sharing:

- Inability to publish data as the consent has not been properly obtained or other issues have not been addressed.
- Rejection when publishing manuscripts. Many journals (e.g. PLOS, BMJ Open Science, Nature Research journals) require supporting data as a condition of publication. If the data cannot be made available then the manuscript may be difficult to publish.
- Research data resulting from publicly funded research are subject to Freedom of Information (FOI) Act challenges. Receiving an FOI request does not mean that there has been a mistake in the data release, however, FOI challenges are much less likely to occur if there has been a properly controlled data release process. FOI challenges should ideally be seen as the last resort for getting hold of the data and not the only way to obtain it.

Regulatory Frameworks

Looking at the wider regulatory framework there are three key pieces of legislation that affect how research data is managed and accessed: the General Data Protection Regulation (GDPR)[2], the 2018 Data Protection Act (DPA)[3] and the 2000 Freedom of Information Act (FOI)[4]. Zosia

[2] https://gdpr-info.eu/
[3] https://www.gov.uk/data-protection
[4] https://www.legislation.gov.uk/ukpga/2000/36/contents

covers a summary of this legislation and some of the legal aspects surrounding data are covered in more depth in Michele Voznik's talk in Section 4.2 of this report.

GDPR / DPA GDPR and the data protection act define expectations and responsibilities around participants **personal data** including special category data, which is personal data that's especially sensitive so extra care must be taken when working with it. These pieces of legislation also set out exemptions for research, where the rules are slightly relaxed if data is being handled for the purpose of research. A key role of these pieces of legislation is to clarify what the requirements are for properly informing research participants of the use and reuse of their data. These requirements are not new to this legislation as the requirement existed in previous versions of the DPA, however, the new legislation contains much clearer guidance on what participants must be informed about and how they must be informed. Zosia also highlights that anonymised data is not subject to GDPR so if you are working with anonymised data in your research then these particular pieces of legislation don't apply.

FOI The FOI is a piece of legislation that covers a broad range of information from public bodies, however, in the context of research it allows the general public to request access to publicly funded research data. Where the risk to the participants is perceived to be low the information is likely to be deemed suitable for release, as has previously been ruled by the information commissioner's office (ICO). An example included by Zosia is the case against Queen Mary University of London (QMUL) where they were faced with a request to release data related to the PACE trial regarding radio exercise therapy in chronic fatigue syndrome. QMUL declined the request on a number of grounds including perceived disclosure risks to the participants. However, the ICO ruled[5] that the data should be released as the disclosure risk was low. This was appealed but the decision was upheld and the data was in fact released.

Having the ability to request publicly funded research data through FOI does not mean that a controlled data release can be disregarded. In fact, Zosia reminds us that the controlled data release is the safer way to release data. However, it is important to be aware that FOI is a route that individuals can take to access research data.

Managing Risk

While we cannot guarantee to completely eliminate all sources of risk to the participants, there are several ways in which we can minimise and manage the risk to participants. In this section Zosia presents several areas in which the risk can be lowered to allow us to make data available in a safe way.

Informed Consent Informed consent is making sure that research participants are informed of any plans that you have for data sharing: what data is going to be shared, how it's going to be made available and who will have access to it. Zosia highlights the distinction between ethical consent for participating in a research study and consent as a legal basis for data processing within GDPR; as these are not the same thing. These different consents could be included in the same form but always ensure that the form is clear. Consent may not always be selected as the legal basis used for your GDPR processing, other applicable bases may include public tasks within the research context. It is extremely important to ensure that your participant information sheets and consent forms are always clear in their meaning, including the legal basis for data processing and what consent you are seeking. This could be seeking consent to share their data or consent to participate in a study, where participation is contingent on their data being shared.

[5] https://ico.org.uk/media/action-weve-taken/decision-notices/2015/1560081/fs_50565190.pdf

Ensuring that your consent forms are as unambiguous as possible is important, especially as the participants are unlikely to have a background in data management. An example of how consent can be confusing was shown by Zosia[6], using two clauses selected from a single consent form:

- I understand that my data will be kept confidential and that only the research team will be able to access information about me.
- I understand that the data may be used by other researchers after the project ends, for work unrelated to this project.

These two clauses seem contradictory with the first clause keeping data confidential and the second clause sharing it. While people from a data management background may infer that the first clause refers to personal administrative data and the second is possibly anonymised data, it is not clear within the form and could easily be interpreted differently by participants from other backgrounds. A participant agreeing to this consent form is unlikely to be certain what exactly they were agreeing to. You should make sure your participants are fully informed and be clear about what will happen with their data at each stage of the research life cycle.

Data Minimisation Another technique that can be applied to data collection and initial processing to manage risk to participants is data minimisation. The key concepts of data minimisation are: limiting the amount of identifiable data collected, and anonymising the data as soon as possible. This reduces the amount of potential disclosed information because if the data is not collected then it can't be disclosed. When thinking about data minimisation you should consider:

- **Minimum variables:** What are the variables that you plan to collect, how can these be reduced down to the minimum amount that you actually require?
- **Level of precision:** If you require age as a variable do you need the exact age, or will age brackets be sufficient? Or if you require a date of birth, would month and year be enough, rather than day, month and year?
- **Masking identity:** To anonymise the data as early as possible, consider whether you can use a pseudonym.

All of these areas try to reduce the amount of identifiable information that you hold as this lowers the risk overall.

Anonymisation The third risk management technique covered by Zosia in this talk was the use of anonymisation. Data anonymisation can be grouped into two types; data alteration techniques and functional anonymisation.

Data alteration techniques include methods such as formal and statistical anonymisation. Formal anonymisation is the removal of direct identifiers in the data set and statistical anonymisation is masking parts of the data that could be used directly, or in combination with other identifiers, to identify a participant in the data set.

Functional Anonymisation on the other hand is not altering the data itself, but rather the environment in which the data is contained or analysed. Measures that can be implemented include: controlled data release, data access agreements to limit what can be done with the data, or release via a data haven. Data havens are special controlled environments that can be much more controlled in their data sharing. Data is held within them and can be analysed but only summary statistics or analysed results are output. The data itself never leaves the haven

[6]Found at 12:28 in the video - https://youtu.be/J9kWkzK83i4

environment. In the case of functional anonymisation Zosia highlights that GDPR/DPA are still likely to apply as the data itself has not necessarily been anonymised.

Conclusion

In conclusion Zosia's sums up some of the key points to consider when thinking about your ethical data management.

Plan early: You need to have a plan before starting. Not only should you plan, but you should plan early. Think about what you want to share, how it will be shared and who you're going to share it with.
Inform participants clearly: These considerations from your plan feed into your participant information and consent forms so that you can ensure that your participants understand exactly what they're agreeing to.
Minimise data collection: Where possible reduce the amount of data you collect, because if you don't have the information in the first place, you can't accidentally release it.
Factor in resources: The processes around data anonymisation and data preparation do have a cost associated with them, make sure you factor this into your overall research plan. This may be just a time cost, or it may be for specialist software. If you don't plan for this you might not be able to release your data at the end of your study.

Resources

In terms of best practice guidelines there are quite a few places to go for support and guidance on these issues. Zosia highlighted a few resources that have been valuable:

- The information commissioner's office provides useful guides to GDPR and the data protection act in quite plain english and also checklists for creating privacy notices.[7]
- The NHS health research authority[8] has some good templates for privacy notices and consent forms.
- The Inter-university Consortium for Political and Social Research (ICPSR) has good recommended language for consent forms to ensure you are clearly conveying your plans for data sharing.[9]

Anonymisation guidance

- UK Anonymisation Network has created an excellent handbook called the Anonymisation Decision Making Framework[10] which goes through a lot of techniques for an organisation and how they can be applied.
- UK Data Service also provides excellent guides on anonymisation.[11]
- Microdata disclosure risk demo from the Computational Privacy Group at Imperial College London[12] provides a website demo of a 2019 privacy study [3]. The study looked at a sample of census microdata and tried to ascertain how many different attributes would be required to identify participants. It found you only needed approximately 15 attributes to identify 99% of participants in this survey. This website allows you to choose which particular attributes you use and see how that affects the disclosure rate.

[7] https://ico.org.uk/for-organisations
[8] https://www.hra.nhs.uk/
[9] https://www.icpsr.umich.edu/web/pages/datamanagement/confidentiality/conf-language.html
[10] https://ukanon.net/framework/
[11] https://www.ukdataservice.ac.uk/manage-data/legal-ethical/anonymisation.aspx
[12] https://cpg.doc.ic.ac.uk/individual-risk/

Anonymisation tools A number of anonymisation tools recommended by Zosia are:

- R package sdcMicro[13]
- ARGUS[14]
- ARX[15]
- QAMyData[16]

The sdcMicro package is aimed at numerical data but along with ARGUS and ARX it allows you to load in data and examine the overall disclosure risk, you can then apply anonymisation techniques and see how that affects the disclosure risk.

QAMyData from the UK data service has a wider scope than just anonymisation, you can look for particular strings that match a direct identifier format, e.g., telephone numbers, credit card numbers or email addresses. In addition, it will also check other areas such as, labels on data variables and missing data. This makes it quite useful in the broader quality checks that should be performed when getting data ready for release.

Questions following the presentation

Q: If data from publicly funded research are under embargo (e.g. say, for 3 years to give the researchers time to publish on the basis of the data) can a FOI request be filed asking for that data to be released before the embargo period is up?
A: *No, one of the grounds for refusing FOI request is intention to publish later on, but you probably have to show that in some written form, by having a practical published or a layout of your publication intentions. If you have intention to do further work on the data and publish data further down the line, then that is a legitimate reason to refuse an FOI request.*

Q: You talked about the dangers of datasets potentially being linked together to make individuals identifiable. One of the advantages of making data available, as say Linked Data so it's in an interoperable form, is that it is easier to link datasets together, are there any additional mitigations you need to take beyond the ones you mentioned when working with linked datasets?
A: *Yes, that is the value of a lot of data, e.g. NHS data, where you have an identifier to link it to other datasets. I think beyond making sure that participants are informed, it would just be a case of storing that linked data in a secure manner, so that the risks of inadvertent disclosure are minimised, and just making it clear what you were doing. You probably wouldn't be able to publish the derived linked datasets themselves, but you would be able to direct people to the source datasets, and if they wanted to recreate that, they could do. That would be the approach that I would recommend.*

Q: You mentioned the UKDS, I'm part of a project where we recently had to fill in an extraordinary amount of paperwork to try to get some postcode based data out of UKDS. It felt like such a long, convoluted process and it was so difficult. We've done it because we really need the data for our research project, but I wonder, is this intentional to put people off or do you think some of those processes should be easier? Or do they need to be that complicated so that you know you've got all the things in place that you need, if you're going to then allow that data out?
A: *That's a very good question, I think there's definitely a certain level of needing to make*

[13] https://CRAN.R-project.org/package=sdcMicro
[14] https://research.cbs.nl/casc/mu.htm
[15] https://arx.deidentifier.org/downloads/
[16] https://www.ukdataservice.ac.uk/about-us/our-rd/qamydata

sure that anyone who wants to use the data, is aware of the risks inherent in it, and they know how to take care of it properly. So there's definitely elements of that to it, but I do think a lot of it could be a bit streamlined. I think there are moves towards streamlining those processes, for example, the new SafePod network[17] *which should be up and running by early 2021. That should be a great help in streamlining access in that way. The idea being that you just have to apply once to get accredited and then you can access a number of datasets from several different data centers in very secure locations. You don't have to jump through the same hoops in terms of getting your own office accredited as a safe location for using data, which can take months and months. So I think there are definite steps to streamline this process, the issue is just that there's a lot of different data centers who have approached the same issues in slightly different ways and it'll take time for that to be harmonised.*

[17] https://esrc.ukri.org/files/funding/funding-opportunities/safepod-brochure/

4.2 Data legislation, personal and non-personal data, ethical issues and protecting your IP rights - Michele Voznick (Pinsent Masons LLP)

Figure 2: Michele Voznik

The full video of Michele's talk can be viewed here: https://youtu.be/w5v5d6r6irs [4]

Michele specialises in privacy, data protection and information law. Her previous work with regulators in the UK and Europe means she brings a unique depth of knowledge and understanding to advising clients on the interpretation and application of this complex area of law. As a UK national expert, she was seconded from the Information Commissioner's Office to the European Commission to work on the GDPR, e-privacy reform and the Digital Single Market. Michele has extensive experience of advising national and multi-national corporations on compliance and data use beyond compliance to optimise their data strategies. She is in strong demand as a speaker, writer and trainer on all aspects of data protection, privacy, Freedom of Information and Environmental Information.

Michele's talk focused more closely on the legislative side of data including details on specific legislation and examination of court decisions. In particular, Michele picked up on some of the issues surrounding protection of data and intellectual property as well as obligations, lawfulness and fairness. The themes within this talk mesh nicely with the talk by Zosia Beckles presented in Section 4.1 of this report. Michele comments that data information law is a continually evolving area with new developments frequently occurring. Some recent developments and guidance for these developments are also covered in this talk alongside a commentary on non-personal information and data ethics.

General Data Protection

Michele begins by talking a little more in depth on the General Data Protection Regulation (GDPR) which was introduced within Zosia's talk. In the UK personal data is defined as something that identifies or is identifiable relating only to a natural living person. However, in some European countries it covers people who have died recently. When undertaking cross-border research you may need to look more closely into the regulations in each country.

There are seven key principles that are the foundation of the GDPR. While these key principles may have been touched on already in some areas, there are the fundamentals to consider when processing personal data.

- **Lawfulness, fairness and transparency** (Article 5.1(a) GDPR) The first thing is it has to be lawful, it has to be fair, and it has to be transparent. Those go together and always get thought about as one group.

- **Purpose limitation** (Article 5.1(b) GDPR) This considers a defined purpose for the data collection and use, and goes back to participant consent; what are they really giving consent for? There needs to be a clearly defined purpose for the use of the data, and if intended to be used for another purpose this must be checked that that's allowed, otherwise it's a breach of the law.
- **Data minimisation** (Article 5.1(c) GDPR) Data minimisation was highlighted in Zosia's talk but is a key theme that will also run throughout this talk. Alongside data minimisation, 'privacy or data protection by design and by default' is also a concept in personal data protection. In this you should think about the data required to get the outputs or research that you're doing. This gives inherent conflict between having the minimum data possible and big data research where researchers want lots of data to see what comes out of it.
- **Accuracy** (Article 5.1(d) GDPR) The next main principle is accuracy, you should be using data that gives you an accurate output or an accurate input, so you should ensure it is fit for purpose.
- **Storage limitation** (Article 5.1(e) GDPR) Also referred to as retention of data, don't keep data for longer than needed and if not needed in a personally identifiable format i.e anonymised data, then anonymise it sooner than later and then it can be kept for longer.
- **Integrity and confidentiality (security)** (Article 5.1(f) GDPR) This is largely covered by saying 'keep data secure'. That means within an organisation and beyond the organisation you don't want a data breach or data being leaked to the public. However, security within an organisation should also be considered, where only those people who need to see the data should be looking at the data.
- **Accountability** (Article 5.2 GDPR) A final key part is accountability and being able to demonstrate that you've met your obligations under the law. As rules for data use become more prevalent and stricter this may well result in more documentation and forms to demonstrate compliance.

Michele concludes this section by reiterating that these principles should form the framework of our data use considerations. If these are always kept in mind, and regularly returned to, then that forms a solid base for responsible data management.

Legal Basis

Within the first principle, lawfulness, fairness and transparency, you must identify the legal grounds for processing of personal data. These are referred to as the 'legal basis' and there are six bases that can be used to process personal data. These are; consent, contract, legal obligation, vital interests, public task and legitimate interests. Michele briefly covers 3 of the bases most likely to be encountered in the research environment.

- **Consent** - Michele highlights that consent has been mentioned already in this session, however, it's one of the bases that is going to be encountered the most, particularly if the research involves going directly to data subjects or individuals. There are some key considerations to remember when using consent for the purposes of the legislation which might be different to the purposes of other your research obligations.

  You have to use clear language to tell participants exactly what you're doing with respect to how you want to use their data. When using the data in more than one project you need to say 'we want to use your data for this specific project and this project that will follow on from it', or stating the specific purpose. For example, 'we want to use your data for three cancer studies and they will be done by these people', you must also tell them

who you are and what you're intending to do. Michele refers back to the examples given in Zosia's talk[18], as these highlighted some of the issues very well.

In addition you must inform them of their right to withdraw consent. This means at any point in time they can say that they don't want their data to be used in your study anymore. This requires documentation to ensure they can do that and should be as simple as them giving their consent in the first place.

- **Legitimate interests** - Within research another legal basis that is likely to be encountered is legitimate interest. Here you have to undertake a balancing exercise between your legitimate interest to be able to do research and the individual's right to privacy. This considers issues around pseudonymisation and sources where data is obtained.
- **Public task** - Another legal basis that might be encountered, depending on how you're doing the research, is public task. This is more likely to apply if the research is government mandated research or required by law.

In addition to a legal basis for processing data, you have extra levels of protection for data that is termed 'special category data'. This is sensitive data, as mentioned before, that deals with race, political opinions, health data, genetic data, biometric data, when it's used to identify an individual's political views or sex life etc. These are areas that are quite intrusive in people's lives and therefore they have an extra level of protection. When dealing with special category data, you need explicit consent, rather than just consent to use the data for research. You have to be much clearer on the data on the data you plan to use, e.g., I want to use your health data for this research, or I want to use information about your ethnicity for this research.

Research Exemptions

Michele then goes on to note that there is however recognition within the data protection legislation world that research is important, so there are a few exemptions relating to the processing and use of data for research purposes.[19] While you can use explicit consent to use sensitive data you also have the research exemption; it says that you can process the sensitive data as long as you have certain appropriate safeguards. These appropriate safeguards are in effect data minimisation, you should only be processing the data that you need and you have safeguards such as pseudonymisation. The research should be done without really making any distinctions about one person, looking at it in an aggregate way and not infringing on people's privacy. If you can do the research with anonymised data then you should be doing the research with anonymised data. Michele also notes that research is quite broadly defined so it covers anything around technology, development, fundamental research, applied research and even privately funded research.

One of the considerations you have when processing personal data is that the data subjects (people) have rights, one of which is transparency, i.e. knowing what people are doing with their data. They also have rights around accessing the data that you hold about them, objecting to the data processing, deleting their data and taking their data somewhere else. There are a number of rights that we all have irrespective of who is processing our data.

Within the DPA in the UK there are some exceptions[20], which are called derogations, such as those enabling use of sensitive data within research. There are also several actions related to data rights that are not required, as long as it is within the research and you've met your other

[18] Shown at 12:28 in Zosia's video - https://youtu.be/J9kWkzK83i4

[19] https://ico.org.uk/for-organisations/guide-to-data-protection/guide-to-the-general-data-protection-regulation-gdpr/exemptions/#ex17

[20] https://www.legislation.gov.uk/ukpga/2018/12/schedule/2/part/6

obligations. For example, if someone says they want all the data that you hold about them, you don't have to do that, you don't have to rectify data if there's a problem with it and you don't have to give them their data back so they can take it somewhere else.

Looking again at the ICO guidance[21], a key point is that if the actions listed in the exceptions e.g. giving someone their data, or correcting data, would prevent or seriously impair your research then you are exempt. However, you must be aware that you cannot output any results from the research in a form that would allow a living person to be identified, you must output the results in a sufficiently aggregated way or in a way that you can't work back to find out who someone is. Michelle links this back to the earlier example of the Netflix datasets and questions on linked data, and highlights that looking toward big data this may become more difficult for researchers to do.

FOI

Michele then returns to discuss the FOI, that was introduced in Zosia's talk, and some of the additional exceptions to FOI requests that apply to research. Michele highlights that the Queen Mary case discussed in Zosia's talk was a really good example of FOI requests and through her work at the ICO she has dealt with a number of these issues. Several additional examples are also mentioned in Michele's presentation.[22]

In addition to the future publication exemption there is also an additional exemption, section 22a, known as the research exemption. This provides an exemption for obtained or derived information if the research program is continuing with a view to publication in the future and there is a prejudice of releasing the data too soon. You can have a controlled data release to other researchers but not making it available to the world at large. However, a public interest balancing test must be done, checking the public interest in the research being carried out and only published when it's ready to go out.

There is an additional exemption that means you do not even have to confirm the existence of information requested through an FOI request, if simply knowing about the information would be prejudicial to the research being carried out. You may just be able to say we can't confirm or deny that we even hold this information so that nothing is revealed.

Decisions around allowing or rejecting the FOI requests and use of the research exemption may often come down to ICO decisions. In particular, there have been quite a large number of FOI requests surrounding the PACE research mentioned earlier, Michele highlights a recent follow on case from the Queen Mary example where they had to release the information. However, there is also an ICO decision relating to King's College and the same research where the research exemption was upheld and the data did not have to be released. So the decisions are based on the particular circumstances and the balancing tests.

Accountability and AI

Accountability is a very important aspect, in fact it is one of the principles of the GDPR that you have to do. It gives you responsibility for complying with data legislation and demonstrating your compliance. Michele touches a little on how you might have to change your accountability considerations when you move towards highly data driven approaches and implementation of techniques like machine learning and AI.

[21] https://ico.org.uk/for-organisations/guide-to-data-protection/guide-to-the-general-data-protection-regulation-gdpr/exemptions/#ex17

[22] The section on FOI begins at 12:10 in Michele's video. https://youtu.be/w5v5d6r6irs

In the UK we take our guidance from the ICO on some of these AI data protection issues, they have produced a selection of guidance relating to accountability and AI. There is the ICO accountability framework[23], which is designed to help organisations demonstrate that they have met their accountability obligations. They have also published ICO guidance on AI[24], which covers what they consider to be best practice for data protection-compliant AI, as well as how they interpret data protection law as it applies to AI systems that process personal data. They recognize there are a lot of benefits to AI, but they also recognize that it poses significant risk to the rights and freedoms of people and compliance challenges for organisations who want to use AI, including researchers.

In their guidance they say that you should embed 'data protection by design, by default' into your culture and your processes. This could be technical issues, but should also be policies and operational issues, which may be difficult with AI systems but needs to be done, and needs to be done at a senior level. The buy-in should be a team effort across the entire organisation not just left to the researchers, data scientists or engineers, because it has to come from the legal compliance angle. The approach adopted should be a risk-based approach, so if you have low-level data that's not very risky to an individual, then your accountability obligation is a bit lower. However, if using high-risk data such as health data, or sensitive data mentioned earlier in these talks, that would be higher risk and you would do a data protection impact assessment, particularly if you're doing large scale processing of sensitive or special category data. You have to set out what you're trying to do, the purpose of it, what the risks are to these people and the steps to mitigate those risks.

There may also be other considerations around data processing, looking at AI and looking at the ICO's guidance. In data protection terms there are controllers and processors. A controller is the person who decides what data will be processed and why, the processor is who deals with the data on their behalf and basically does what they say. In an AI algorithm relationship this becomes a bit more complicated and that's something that will have to be developed in time, but it's is an area that should be thought about as you enter into relationships with people to undertake processing for you.

Court Decisions

Michele touches on the Schrems decisions briefly with examples from French data protection decisions. The Schrems decisions, Cases C-362/14 and C-311/18 of the European Court of Justice, relate to data transfer mechanisms between countries, in particular coming in and out of the European Union. When transferring data abroad, to the European Union or any other country, there are certain restrictions and safeguards that need to be in place, however, the recent court decision ruled that the previous safeguards were not actually adequate protection and additional mechanisms were required. This is a crucial consideration if you are looking at using a cloud provider or data processor abroad and these issues should be discussed with your legal team to ensure that your contracts are fit for purpose. Michele notes that new guidelines have only very recently been published, so this is an area that lawyers are currently trying to see what this means in practice for data transfer.

As an example of how serious the Schrems decision was, Michele highlights a decision by the

[23] https://ico.org.uk/for-organisations/accountability-framework/

[24] https://ico.org.uk/for-organisations/guide-to-data-protection/key-data-protection-themes/guidance-on-ai-and-data-protection/
https://ico.org.uk/for-organisations/guide-to-data-protection/key-data-protection-themes/explaining-decisions-made-with-artificial-intelligence/

French data protection authority (CNIL), the French equivalent of the ICO. CNIL ordered the French health data hub to stop using Microsoft as a cloud provider for their COVID health research hub because they said that it didn't comply with the Schrems decision. Now that decision was referred to the French courts and they said they disagreed with the regulator. This is a very important area and measures need to be put into place but they're not going to just stop the flow of data. This is likely to continue to be a source of conflict as everyone adjusts to what the new rules are. There may be decisions by regulators that get overturned by the courts, or courts which take a more protective approach to people's rights.

Michele highlights a privacy non-profit organisation called None of your business[25] (noyb) by Maximilian Schrems, the man behind the complaint that kickstarted the Schrems decision that brought down two internationally agreed data transfer treaties. Noyb is very active in making sure that people's personal data is protected and processed appropriately. They've made over 100 complaints across Europe about businesses and other organisations processing personal data. If you're interested in data legislation enforcement and whether it may affect you they also have a GDPR hub website[26] where you can see decisions from various courts relating to data privacy. There is also a GDPR enforcement tracker[27] which reports fines imposed under the EU GDPR.

Non-personal data

Michele then moves to talk about non-personal data and how this can differ from the personal data that has already been spoken about. When you're dealing with non-personal data the GDPR doesn't apply, if you can genuinely anonymise your data then that is outside the realms of the GDPR. This may be the case with industrial data, anonymised personal data, certain statistical data, these can all be non-personal data and the legislation mentioned above does not apply, with the exception of the freedom of information act, which may still be applicable if you are a public body or university.

There is also the case of mixed data, when you have non-personal data and personal data all mixed together, how can this be treated? The general rule is if you can genuinely separate out the personal and non-personal data then you can process the non-personal data without worrying about the GDPR and related obligations. However, if the data is inextricably linked then you have to apply the GDPR to the whole dataset and treat it all as personal data.

Michele highlights that the European Commission is supportive of research and developing data, so they're looking at making sure that data (non-personal data) can flow freely across the European Union. There are a lot of policies and initiatives in place to ensure that there's free movement of data and good self-regulation using cloud security. They are also setting up a European cloud for research which may be important to universities and researchers in the future.

Data Strategies

Data strategy is quite a hot topic at the moment, the European Union has a data strategy and the United Kingdom also has its own national data strategy and if you're dealing with other European countries they'll probably each have their own data strategy as well. However, researchers are usually well considered within these strategies. The data strategy is in effect trying to make the EU or the UK more competitive globally and a leader in a data-driven

[25] https://noyb.eu/en
[26] https://gdprhub.eu/
[27] https://www.enforcementtracker.com/

society. There are many developments that are coming in this area as many are still in the early stages. Some developments may come in 2021, as many things that were being worked on in 2020 have been pushed back by COVID. However, COVID has also highlighted the need for better access to data especially for research purposes.

Michele highlights that while Brexit will impact these areas it is still useful to have knowledge of the European strategies. Within the EU, their strategy is looking at a single market where data can flow easily for the benefit of many, including researchers. They are trying to set up clear rules for access to data and reuse of data, and looking at different issues including investing in infrastructures, European cloud capacity and interoperable data spaces. Examples of projects include the European health data space[28] and the European Open Science Cloud.[29]

Looking at the UK national data strategy, it's very similar to the European one. The main driver within it is to find ways that data can be used responsibly, lawfully, securely, fairly, ethically, sustainably and accountably for supporting innovation and research. New legislation is not planned, but they are looking at ways that non-legislative measures can assist. These include initiatives like NHSx[30] which is the data strategy for health and social care, and opening up datasets for that.

Data Ethics

Michele briefly touched on data ethics, which was the subject of our breakout discussion sessions captured in Section 5 - Interactive Breakout. Ethical use of data is a key area of consideration and Michele highlights a few ways in which data ethics is beginning to appear in the compliance area and organisations that are examining it.

- At a European level the European data protection supervisor (EDPS) has a five year program looking at data ethics and how that works with data compliance, because as part of getting your compliance for personal use of data, ethics comes into the fairness element of it.
- The information commissioner's office in the UK have for the first time appointed a data ethics advisor[31], and are now looking at what data ethics means for the use of personal data in the UK and how that fits with meeting compliance obligations.
- Data & Society[32], from the US, have done a lot of research around big data and technical research; highlighting some of the ethical issues brought around by that. In particular they have published a working paper around supporting ethical data research.[33]

Intellectual Property

Michele discusses intellectual property considerations and includes some top tips on areas that should be thought about with respect to intellectual property when starting out on a research project. Michele notes that she is not an intellectual property lawyer, but often works closely with them on data issues, and these are some important areas to think about.

- **Collaboration:** If you are collaborating on your research it's important to consider who's going to own, manage and exploit any IP that you already own, or that might be created.

[28] https://ec.europa.eu/health/ehealth/dataspace_en
[29] https://www.eosc.eu/
[30] https://www.nhsx.nhs.uk/
[31] https://ico.org.uk/about-the-ico/news-and-events/blog-data-ethics-and-the-digital-economy/
[32] https://datasociety.net/
[33] https://datasociety.net/library/supporting-ethical-data-research-an-exploratory-study-of-emerging-issues-in-big-data-and-technical-research/

You need to think about contracts in advance of the project, because you need to consider these issues before you start signing contracts, either with funders or with collaborators.

- **Scope of the project:** What is the scope of the project? This also comes as a consideration under the GDPR when talking about the purpose, you need to know what the scope of your project is. You need to know what potential IP rights are going to come out of it.
- **Ownership:** Who is going to own the IP? You might want to consider licensing or royalties. It may not be one party's desire to keep all the IP, because once you've got IP you might have to maintain and enforce it, which might not be something you want if you're a small research group. If you have more than two people you need to know where the chain of responsibility is and who is responsible for what aspects of it.
- **Use and Exploitation:** You may also need to consider background IP, if using any background IP. To what extent is this needed for the research that you're doing and how can you ensure that it's licensed appropriately for the other parties to use.
- **Infringement and Enforcement:** Expanding on who owns the IP, who is going to be responsible for dealing with infringement of it or enforcing it? Do you need the consent of everyone, will someone specific be in charge of that, who would fund any litigation? You need to consider it as litigation can be expensive and you might not want to go down the litigation route.
- **Confidentiality:** A very important topic that considers data protection, keeping data secure, the controlled data release program, freedom of information and potentially other issues around IP. You don't want to disclose it, because that could affect the granting of a patent or creation of a design right.
- **Exiting:** All projects come to an end at some point. At the end of a project you don't want to have all these awkward conversations that you should be thinking about at the beginning of the project.

Michele rounds up her talk by talking about data ownership. Data, whether that's non-personal data or personal data nobody owns it. People may have rights over it, obligations to it, be able to exercise obligations over it, but there is no ownership. In the considerations you have to look at contracts, GDPR, FOI, IP rights and Michele refers to it as a cornucopia of data legislation to consider. As a final point Michele highlights that these are all things to think about at the beginning of a project rather than the end of a project.

Further Reading

In addition to the resources recommended by Michele in the footnotes of the talk, here are a few additional resources recommended by Michele.

- The Ethical Alogrithm by Michael Kearns & Aaron Roth [5]
- Guidance from the ICO - https://ico.org.uk/

Questions following the presentation

Q: If you collect data with consent and then anonymise it and then a participant withdraws, does this affect anonymised data and how can you withdraw it without it being identifiable?
A: *The short answer is, once you've anonymised it, it's no longer personal data. So that withdrawal of consent doesn't affect it. If you still hold the identifiable data, then you have to stop processing it. Any processing you did before the withdrawal was made stays lawful. You just wouldn't use any of that identifiable data going forward. The other aspect is, under the GDPR you don't have to collect additional data to allow someone to exercise their right. So if you don't know who it is, you don't have to go and ask for more data just to be able to deal with that.*

Q: My understanding was that you could give limited rights to withdraw data if after anonymisation it would be impossible to withdraw and/or detrimental to the project?
A: *So again after anonymisation it is no longer personal data, so for GDPR purposes you don't have to worry about that because then you're outside the realm of the GDPR. But with withdrawal, if you're talking about deleting the data, which is slightly different whether you're using consent or using another legal basis, then there are certain elements around the right to be forgotten or erasure. If it's potentially detrimental to your research project, then no you don't have to delete the data, but you've got to make sure you fit your other safeguards in there.*

Q: What has everyone done about voice recognition, automatic transcription software and transcription services? In particular the ethics/ legalities of students using AI driven transcription in another country. Is this identifiable information? (Based in a South African University)
A: *The voice files recorded for the purposes of transcription are identifiable to a certain extent. My picture, my voice etc. is something that's identifiable if you're using it for identity purposes. If the people have given a consent or they've signed up to the research project it is not such an issue to use a transcription.*
With respect to sending the data from South Africa internationally, I'm not a South African lawyer, I know you have data protection laws, but I'm not 100% sure that you have the same international rules that we have in Europe. In Europe we have very strict rules about international transfers. You may have to have a contract, so I would check with your legal Department of the University. They can access this information to see that the service is fit for purpose; there are the security elements, that it is accurate in what it does. If the University has a contract for the provision of that service, whether it's free or not free, that it has safeguards particularly on confidentiality and security. These would be my top tips on that one. But unfortunately I can't comment too much on the specifics as I am a UK lawyer and I'm not qualified to speak on South African law. It's something I would definitely take up with your legal Department to ask those questions.

Q: A question around confidentiality, and this goes to the larger issue. How confidential is an AI?
A: *Well, that's a big question. So when you're looking at confidentiality, it's fine to use algorithms. What you want to do if we're just purely talking about confidentiality and security is to make sure that whoever is processing the data through the algorithm, whoever is running it, that their systems are secure. Making sure your systems are secure and how you get it to them is a secure method of transfer. Those are the security issues we're looking at. Then you come into the accuracy and the other issues, but I would say on the security, that's the one you're looking at for the AI.*

Q: No one can own research data? So, for non-personal data, then, is it really only about access rights to data and IP?
A: *For the most part yes. You can have database rights and other rights that do exist around data, you can have your trade secrets, but you have to make sure those rights do fit the data. Data in itself is not something physically that can be owned, so it is about rights over it. Basically that is the view about it and we've never found a case where data can be owned because it's not a physical thing and the courts have pretty much decided that. It's about rights and obligations and how you protect your data and how you can utilize your data as an asset.*

Q: If a researcher publishes a dataset as 'open data' under a CC0 license, i.e. so anyone can use it freely for any purpose, and then someone else comes along and reuses that open dataset in combination with another dataset resulting in non-GDPR compliance, can the researcher or organisation who released the original open data set be held liable in any way?

A: *If you send out the open data under license, you've lawfully released that data under that license. Probably it was in an anonymised or pseudonymised format, so that it wasn't personal data. If someone comes along and reuses that data, and actually then puts it with something else and it becomes identifiable, they're the ones who then have the data protection obligations for that data. The other issue in that one is, if data has been anonymised in the UK, it's actually an offence to re-identify it.*

Q: To what extent, presuming we have the administrative and technical safeguards, can universities archive identifiable research data for general research re-use by 'other authorised researchers outside the institution not involved in the original study'? For example, the University of Bristol controlled data Data Access Agreement[34] seems to indicate that in some circumstances identifiable information might be shared outside Bristol under a data access agreement which makes the recipient a data controller. I think the legal basis for sharing personal data in such a way, if feasible, would have to be consent, which makes it in principle revocable by the data subject at a future date. And how do you frame the original participants' consent, since it is for a general research purpose, rather than a specific identified project? (Asked to both Michele and Zosia)

A (Zosia): *So I think in the case of Bristol, I'm not sure whether or not the data access agreement actually makes the applicant a data controller. I think they would still be a processor because ultimate responsibility for the data would still reside with the University of Bristol, although I haven't looked at the text of the agreement in fine detail recently.*

The more general point about can we reuse identifiable data? I don't think there's any doubt that GDPR makes provisions for that in research and statistical purposes. That's somewhat one of the reasons why we (University of Bristol) have moved towards recommending that researchers use public task as the legal basis for data processing in a research context, so that it's clearer that we are separating that ethical consent to participate in the study from GDPR consent for data processing. Then because we're using public task and the additional research exemptions you don't have that issue of participants being able to withdraw their data further on down the line when the dataset is still identifiable, so that solves some of those problems. It's a tricky area, but I think certainly as long as you are being careful with how you're storing the data and how you're managing access to it, I don't think inherently there's a problem with using identifiable data for subsequent research purposes, provided those purposes are in line with the way it was originally collected and the participants were aware of what could be done with the data later on.

A (Michele): *It can't be incompatible, you need to tell your participants and you need to have your accountability documents in place. But it's not not possible. Although, you do have to make sure that you dot your i's and cross your t's in that respect.*

[34] https://www.bristol.ac.uk/staff/researchers/data/accessing-research-data/

4.3 Practical Ethics for Data Science and Algorithm Design - Tessa Derbyshire (Patterns)

https://orcid.org/0000-0001-5433-5795

Figure 3: Tessa Derbyshire

The full video of Tessa's talk can be viewed here: https://youtu.be/jEFu1ykVI_I [6]

Tessa has a philosophy background, and her academic work, within a global group, focuses on the computer/human interface, particularly in relation to explainability, interpretability, and trust. She has expertise across the fields of technical validity, privacy legislation, and ethical practice, particularly in relation to the human impact of data.

Tessa begins her talk with an anecdote about the title. The title she was originally working with was "Ethical Algorithm Design and Ethical use of Data", but this was discarded due to the nebulous nature of ethics, with Tessa remarking that speakers on this topic often tend to presume that their audience are all well versed in ethics, when in reality people who aren't specifically working in the field might well have an awareness of ethics, but aren't necessarily immediately looking to interact with the topic. Further, there can be a tendency to use obfuscating language when discussing the complex dimensions of ethics, which can both be off putting and can detract from the matter at hand. Tessa suggests that we need to step away from the hallowed halls of classical ethics and formulate something more tangible that is grounded in how people actually behave.

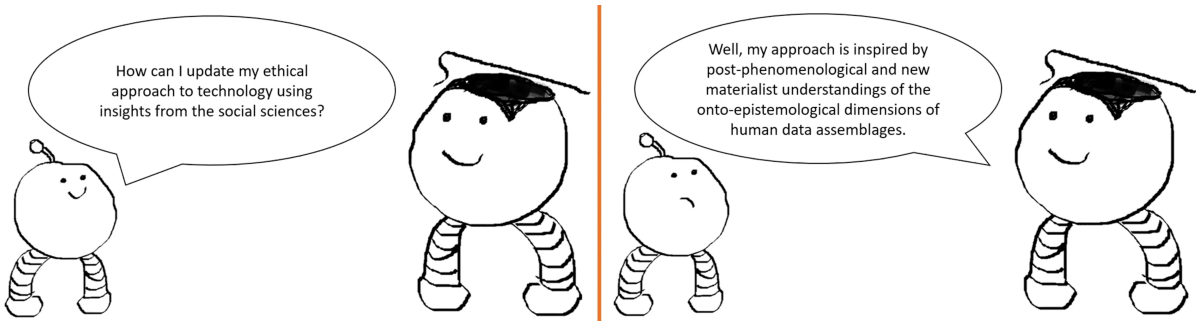

Figure 4: Cartoon from Tessa's talk depicting the use of obfuscating language in discussions about ethics

Therefore, Tessa's talk is focused on how we can talk about practical ethics, with three main facets:

- Ethics as scientific rigour
- Ethics as acting within the law
- Ethics as societal impact

These aspects were picked because of their applicability to data scientists across different domains, and due to the fact that each is covered by well established best practices. The main message of this talk was to emphasise that these structures already exist, and whilst new developments and legislation can change how we handle challenges in ethics and ethical data science, ultimately these pre-existing structures remain in place to be utilised. Tessa explains the existing practices in each topic and demonstrates how they can be applied with respect to ethics.

Ethics as Scientific Rigour[35]

It is widely accepted that researchers have a responsibility to plan and publish research which is scientifically rigorous, reliable and reproducible [7]. The scientific community benefits from this practice when it is highlighted as a commitment in published work because it provides a coherent body of knowledge for future research. This applies at each stage of the research life cycle: research design, data collection and preprocessing, analysis methods, interpretation and theorising. This is true irrespective of the domain that this research takes place in, and as such it is imperative to plan out these aspects as early as possible. There are too many situations where ethical statements are tacked on post project, or a smattering of ethical citations are added at the end of a paper. This is not how research should be conducted, ethics needs to be considered right at the beginning of the research. Tessa details the different stages of the research life cycle:

- **Research design:** The whole purpose of research design is to encompass a clear research question and a description of the problem or opportunity that motivates the work with reference to the existing literature in the relevant field.
- **Data collection & Preprocessing:** This covers the requirement to clearly outline the methods used in collecting and preprocessing data, including weighing the advantages and limitations of the collection process, and the management of the collective data in adherence to principles such as FAIR. ensuring that data is Findable, Accessible, Interoperable and Reusable.
- **Analysis methods:** This refers to the techniques used to extract insights from the data and the expectation that the choices made are justified with reference to the research questions, existing literature and relevant benchmarks or baselines.
- **Interpretation & Theorising:** This represents the requirement upon researchers to ensure that all claims they make are supported; highlighting clearly which features are being used to generate insights, recommendations or predictions.

Tessa argues that these points essentially amount to a set of ethics to be followed with respect to conducting data science research. The result of adhering to the principles of scientific rigour is a body of scientific knowledge which can act as a foundation for advances in technology which can change the way we live our lives. This holds true irrespective of whether you are developing a facial recognition algorithm [8] or a protein folding algorithm [9] or a novel renewable energy technology [10]. The infrastructure for scientifically rigorous research is well embedded across domains, in pursuit of a practical ethics Tessa recommends that it can and should be used as a foundation.

[35]This section starts at 03:31 in Tessa's video: https://youtu.be/jEFu1ykVI_I

Ethics as acting within the Law[36]

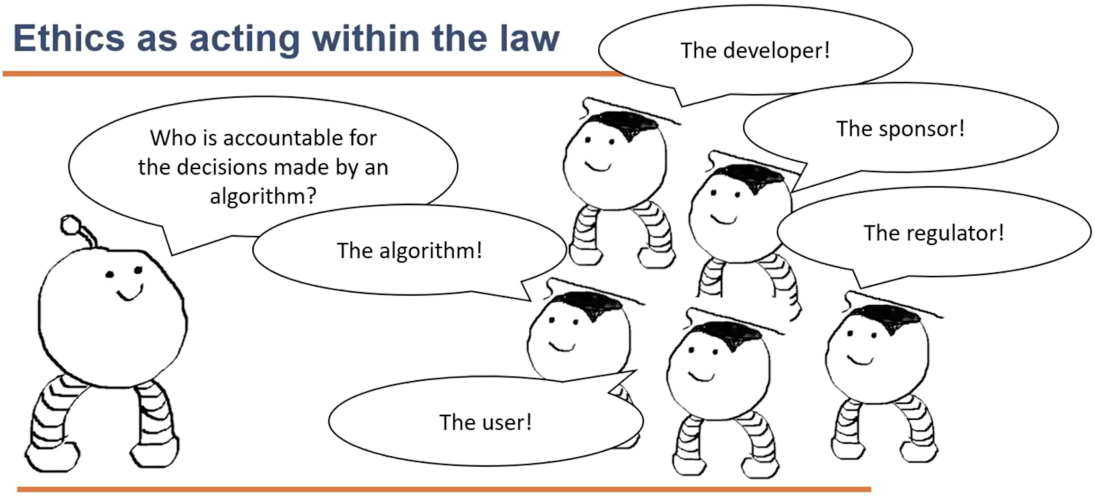

Figure 5: Cartoon from Tessa's talk about ethics legislation

Tessa begins by stressing the importance of accountability with respect to ethics from a lawful perspective. At present there are a multitude of papers and publications discussing the key question about who is (and should be) truly accountable for automated decision making [11,12]. There are publications suggesting options ranging from the algorithm itself, the user, the developer, the regulator, the sponsor, and a variety of other options [13, 14]. Further, there are also papers discussing the questions raised by this dilemma, such as how do we attribute responsibility? Particularly if the decisions are made by a system that cannot easily explain its decision making process. These are extremely complex questions and the volume of scholarly and governance work on the subject reflects this.

Tessa highlights relevant work conducted by Professor Joanna J. Bryson[37] from the Hertie School of Governance as part of the Oxford Handbook of Ethics of AI [15,16]. Joanna's work deals with questions surrounding the explainability of a given process. Joanna holds that we do not need to completely understand how a machine learning algorithm works to regulate automated decision making, any more than we need to completely understand the physics of torque to regulate bicycle riding in traffic [17]. In this analogy there are complex roles and responsibilities played by all stakeholders; the cyclists, the bike manufacturer, the merchant, nearby motorists and pedestrians, the local council, the legislator, etc. Similarly there is a complex web of responsibility encompassing both the development and deployment of algorithmic systems, and data science itself more broadly. As these issues reach applications they increase in complexity, but can be regulated without needing complete transparency of the system itself.

There have been many discussions around whether AI in particular, and general AI as a focus, can be accountable in its own right. [18] Joanna states that "The issue should not be embedding our intended (or asserted) values in our machines, but rather ensuring our machines allow firstly the expression of the mutable intentions of their human operators, and secondly transparency for the accountability of those intentions, in order to ensure or at least govern the operator's morality" [15]. In this analogy, we want to be able to assess that the manufacturer of the bike did not intend for an accident to occur in traffic and they did their due diligence in averting the outcome. Similar points can be made regarding the cyclist, in that they didn't deliberately

[36]This section starts at 06:22 in Tessa's video: https://youtu.be/jEFu1ykVI_I
[37]https://www.hertie-school.org/en/who-we-are/profile/person/bryson

cause an accident and they were following the laws that were set. We cannot reasonably claim to ensure the morality of either the bike manufacturers or the cyclist but we can govern it based on the outcome of their actions.

Tessa highlights that through this analogy it can be demonstrated that there are extensive existing infrastructures such as independent audits, requirements for due diligence and recourse, and the governance process which can act as a solid foundation for practical ethics. Moving forward, there will be a requirement for legislation that specifically focuses on the opportunities and pitfalls of these evolving technologies, building on the General Data Protection Regulation (GDPR)[38], Freedom of Information (FOI)[39] and the Data Protection Act (DPA)[40]. However, the core principles of responsibility remain the same, this is a technology that humans have built and humans are responsible for and every day we take these responsibilities into consideration. Tessa notes that she has completed many office training hours on what it means to behave ethically within your company and these issues are very similarly applied in this context.

Ethics as Societal Impact[41]

This is a more fuzzy area, as impact and response can depend on many factors including location, time and culture. Tessa cites a recent example of the Student protests[42] that took place this year about the algorithm that assigned the exam grades due to exams not taking place during the pandemic. There were many concerns about the viability and accuracy of these grades, however when the students protested, they weren't protesting using words such as "Fairness, Accountability, Transparency" they were protesting against feelings of persecution, classism, and lost university places. That demonstrated real tangible impacts that can happen as a result of decisions that are made about data science and algorithmic processes. Tessa argues that researchers and developers are responsible for considering their work in a social context despite the shifting goalposts that characterise discussions in these areas.

With respect to identifying guidelines that exist in this area, a vast range of organisations across different domains in different locations have been producing ethical guidelines. As noted by Michele in the previous talk, the EU are working on this[43], the UK is developing a National Data Strategy[44] and the United Nations is working to develop guidelines for ethical AI[45] and these are all currently in discussion. Further, many conferences (e.g. The Neural Information Processing Systems Conference (NeurIPS)[46]) require every paper to include a social impact statement. Three major scientific journals (e.g. Nature Medicine[47], BMJ[48] and the Lancet[49]) have announced that they are working to develop domain specific guidelines that govern the publication of research in their area, and to extend the CONSORT[50] and SPIRIT[51] guidelines such that the medical legislation covers the use of AI [19].

[38] https://gdpr-info.eu/
[39] https://www.legislation.gov.uk/ukpga/2000/36/contents
[40] https://www.gov.uk/data-protection
[41] This section starts at 10:20 in Tessa's video: https://youtu.be/jEFu1ykVI_I
[42] https://www.technologyreview.com/2020/08/20/1007502/uk-exam-algorithm-cant-fix-broken-system/
[43] https://ec.europa.eu/digital-single-market/en/news/ethics-guidelines-trustworthy-ai
[44] https://www.gov.uk/guidance/national-data-strategy
[45] https://www.itu.int/dms_pub/itu-s/opb/gen/S-GEN-UNACT-2019-1-PDF-E.pdf
[46] https://nips.cc/
[47] https://www.nature.com/nm/
[48] https://www.bmj.com/
[49] https://www.thelancet.com/
[50] http://www.consort-statement.org/
[51] https://www.spirit-statement.org/

Some applications have a much clearer social impact than others. For example, the use of AI in healthcare has a clear and direct impact on the individuals involved. However, it is worth noting that not all data science is AI, but whilst AI is the current flavour of the month, other methods still need to be considered with respect to their impact. All data can and should be placed in the context of its lifecycle from collection through use, and in terms of practical ethics or ethics in practice. Tessa highly recommends seeking out the groups developing ethical guidelines applicable to your research area.

The talk concludes with a quote from Deborah Raji from the AI Now Institute [20] "There's no such thing as flattering stereotype, only flattening stereotypes. The strong Black woman, the math genius Asian. It is an inherent insult to crush full and complex human beings into the one or two dimensions one chooses to see". On the internet, the world is flat. For reality to be computable it must lose its dimensions. The definition of dynamic and complicated objects including and especially humans collapses into inputs for operations and analysis. This results in an inherent tension between minimisation and privacy. If you are minimising your data within the law for the purpose of retaining privacy, are you minimising it too far? Is that data fit for purpose in understanding the nuances of people? Ultimately data is reductionist, and whether these reductions leave you with something workable is a really important question. Tessa stresses that the data you leave out needs to be considered as carefully as the data you choose to include.

Further Reading

Tessa recommended the following resources for further reading:

- The work of Professor Joanna Bryson[52] on Explainability
- The work of Dr Joy Buolamwini[53] on Algorithmic Justice
- The work of Dr Timnit Gebru[54] on Algorithmic Bias
- The work of Dr Inioluwa Deborah Raji[55] on Algorithmic Accountability
- The work of Abeba Birhane[56] on Cognitive Science

Questions following the presentation

Q: About some of the things that are coming out, for example some of the frameworks or decisions about things going forward - is there anything that really excites you or really worries you about some of the future developments of ethics within the AI sphere?
A: *I have some concern about the new regulations, particularly the EU guidance on Ethical AI. It is quite focused on the opportunities (e.g. innovation and business aspects) and doesn't tend to highlight the threats or potential risks. It also doesn't seem to address the underlying societal challenges that really drive the bias that we see in the algorithmic world. It would be fantastic to see that balance change a little bit as the new legislation comes in.*

Q: You mention FAIR in relation to ethics as scientific rigour. Do you think there is a risk that FAIR may move from the ethics area into the compliance area? I'm thinking about how Michele mentioned the creep of ethics towards compliance.
A: *I think there is creep between the areas, but I don't know that I would describe it as creep. Compliance comes from ethics and ethics come from compliance and the interlinked nature of*

[52] https://scholar.google.co.uk/citations?user=QOU1RTUAAAAJ&hl=en
[53] https://scholar.google.com/citations?user=6SG9440AAAAJ&hl=en
[54] https://scholar.google.com/citations?user=lemnAcwAAAAJ&hl=en
[55] https://scholar.google.com/citations?user=pzw1-J4AAAAJ&hl=en
[56] https://scholar.google.com/citations?user=D1ApV_YAAAAJ&hl=en

these things means that both Michele and I as you've noticed, despite coming from different perspectives, have talked about similar things. I think that FAIR does a reasonably good job of setting clear principles that researchers in particular find easy to adhere to and with the mess of available principles, those clear structures are really needed. Whilst FAIR may move towards compliance, if it helps support the ethical structures that's the way of the world unfortunately.

Q: Do you think societal impact statements have a place in non-human related research, e.g. industry or engineering?
A: *Yes, I do, because as I said I think that data should always be considered in context. Absolutely in general they are aimed at researchers using human data, but I think taking the time, particularly at the beginning of the research phase to think about what the impact might be and incorporating that through only makes your publication or your research stronger because it has those considerations of wider context of the data.*

5 Interactive Breakout

Following on from the talks an interactive breakout discussion was held. The Moral-IT cards[57], developed by Dr Lachlan Urquhart (University of Edinburgh) and Dr Peter Craigon (Horizon Digital Economy Research Institute, University of Nottingham) [21], are used as a tool to prompt reflection on the legal, ethical, technical and social implications of new information technologies. In our interactive session, we used the Moral-IT cards to prompt discussion and thought on the ethical implications that a user's research can have, and how they can consider these angles in their experiment and research design. The session was introduced by Dr Peter Craigon and discussion groups were facilitated by members of our organising team. The Moral-IT cards are licensed under a Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International License[58].

5.1 Introduction

This interactive activity focuses on using the Moral-IT Cards, a tool for ethics by design, a practice that reiterates the messages that have been repeated in each talk at this event. Ethics is something that needs to be considered and embedded from the start, rather than thought of retrospectively. Peter began introducing this activity by explaining the nature of ethics by design:

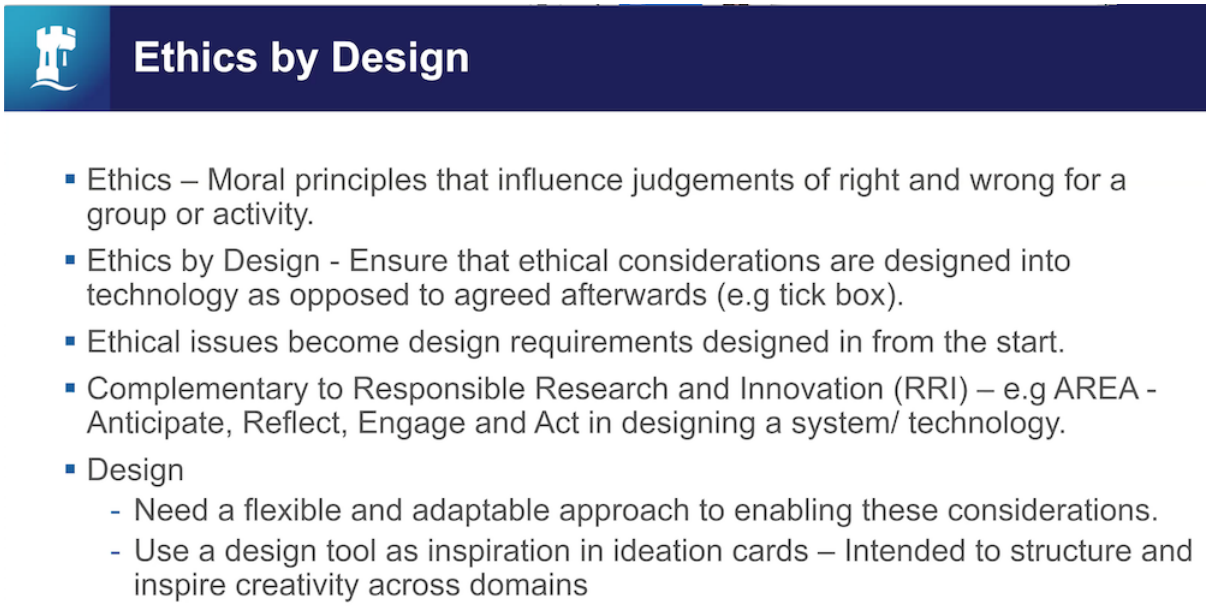

Figure 6: A slide from Peter's presentation explaining ethics by design

Before embarking on a research project or the advent of a new technology, It is highly advisable for researchers and developers to anticipate the effects of what they are doing. They should engage with potential stakeholders and users, and those who might have an interest in it to gain a better understanding of its potential impacts, and use this knowledge to make improvements. Ethics by design encourages researchers to put this level of forethought into their work, and these Moral-IT cards have been developed as a tool to aid with this. The Moral-IT cards are a deck of cards, similar to a set of playing cards. They have been divided into four suits: privacy, ethics, law and security. There are also narrative cards which help structure discussion.

[57] https://lachlansresearch.com/the-moral-it-legal-it-decks/
[58] http://creativecommons.org/licenses/by-nc-sa/4.0/

Figure 7: Examples of the Moral IT Cards taken from the Moral-IT Website[59]

As shown above in Figure 7, each card has the following structure: the title, an image to illustrate, provoke and provide an alternative perspective (perhaps on the principle itself), and an open question which is intended to encourage engagement, which needs to be answered and cannot be dismissed or answered with a simple yes or no. These cards were used in this breakout session to prompt discussion around a proposed piece of technology.

We identified four main topics to be discussed in the breakout groups: data management, data collection, data analysis and data sharing. Each of these were discussed using the cards as prompts to consider different ethical issues associated with each of these topics. The following sections summarise the discussions that were had around each topic, noting the key discussion points and highlighting any specific cards that were raised as relevant. For the purposes of this interactive activity, the Moral-IT deck was split up into selections of relevant (but overlapping) cards that were most pertinent to each topic. The PDF file[60] illustrating the groupings selected for this discussion is also provided for reference.

5.2 Data Management

Data management is the entire process of managing, storing, organising, and maintaining the data that has been collected as part of your research. Most research projects will require a data management plan to state how this will be carried out throughout the project. These plans are extremely important and should be designed at the very beginning of your projects.

The cards that were picked for this section were:

- Trustworthiness
- Sustainability & E-Waste

[59] https://lachlansresearch.com/the-moral-it-legal-it-decks/

[60] https://www.ai3sd.org/wp-content/uploads/sites/374/2021/01/moral_it_deck_FI2NI.pdf Also available as part of the additional files at: http://dx.doi.org/10.5258/SOTON/P0034

Trustworthiness

Trustworthiness was picked because it is hard to know how much trust you can put in the people that are collecting your data, in the people that are storing your data, in the systems that are storing your data, and even in yourself that you actually know how that data is going to be used and stored.

Figure 8: The Trustworthiness Moral-IT Card

Considering this card also raised some further discussion points:

- **Lack of awareness of perspectives:** Metadata is really important, but some will consider that metadata isn't for them, it's for others. Whereas in reality metadata is useful to both the people working on the project and others using the data post project.
- **Variance in access considerations:** There are different perceptions around different types of storage. There is an assumption that physical records locked in a room are more secure than digital records stored on a server. However, this isn't always the case, it really depends on the individual circumstances.
- **Ask questions early:** Asking questions at the beginning is really important. Sometimes this can be challenging at the beginning, particularly if everyone is having fun with the project.
- **Some institutions/people hold implicit trust:** There might automatically be more trust in a University collecting data than a company, due to the perception of what each institution is likely to do with it and how confident users are in them. There is clearly a high level of trust by many in Google and Amazon for storing their data on their cloud services, based on the sheer amount of usage. Equally, there can be a tendency to trust people you know; familiar teams can take things for granted after a while.
- **Reputation impacts trustworthiness:** Trust tends to be based more on reputation than personal knowledge, and bigger companies can find it easier to gain this reputation. E.g. Amazon/Google. Smaller companies have to put more effort into proving that they can be trusted with your data, for example by gaining accreditations for their services.
- **Location is key:** Researchers often prefer to store their data on campus rather than sending it off to a Cloud storage provider e.g., Dropbox where they don't know which country their data will be stored in.

- **Proxies for trust:** There can be proxies for trust. For example, when considering whether to trust a data repository, you might look at who is funding it. There might be requirements to deposit data in certain repositories depending on your funding, but this in itself builds trust as these become well used repositories associated with the funders. Further, if there is the requirement to use a specific repository, there may be less of a desire to scope out potential other repositories to use instead.
- **Legacy data poses an issue:** Ten years ago, many of these issues weren't even being considered. The main concern was about having the data required to conduct the analysis rather than how and where it was stored.
- **More advanced data management plans:** The research data alliance is working towards more advanced versions of data management plans. Both with respect to active data management plans that are working documents that evolve as the project evolves, and for machine actionable research data management plans that enable some of the sections of these plans to be automatically filled in where data is being collected automatically by a machine [22].

Sustainability & E-Waste

Sustainability & E-Waste was picked because this is an area that can be significantly affected by the requirement to store a large amount of data, but is often overlooked. For example, storing large quantities of data can lead to the need for increased air conditioning to keep the data centre cool enough, which can in turn lead to a significant increase in the power required to keep the centre running.

Figure 9: The Sustainability & E-Waste Moral-IT Card

Considering this card also raised some further discussion points:

- **Consider Water Cooled Data Centres:** One way to reduce power usage is to use a local, purpose-built data centre for research data, and use direct water cooling to regulate the temperature. This blog[61] provides further details.

[61] https://blog.bham.ac.uk/bear/2018/07/11/the-launch-of-our-new-research-data-centre/

- **Consider locations for data centres:** It could also be worth having data centres in sensible locations, such as Iceland because it is very cold there and so the cooling will be less of an issue.
- **Start early:** When researchers consult with experts, they do this too late, such as considering sustainability aspects after doing the research. Researchers need to get ahead of this and do the work at the beginning.

5.3 Data Collection

Data collection is the process of collecting and measuring information on specific variables that have been identified as relevant to your research questions. A researcher can evaluate their hypothesis on the basis of the data they collect. Data collection is a vital step for research in any field, although the methods and approaches can differ depending on the domain and the required information. Data collection, processing and analysis methods should all be described up front in the data management plan.

The cards that were picked for this section were:

- Accountability
- Confidentiality

Accountability

Accountability was picked because it is becoming very common for lots of data to be collected by machines, which poses several questions about the levels of accountability, and who is actually accountable.

Figure 10: The Accountability Moral-IT Card

Considering this card also raised some further discussion points:

- **Data policy is a key tool for accountability:** Rules around ensuring that researchers who do not comply with data policies (e.g. data management plans, formal research data management activity) would be denied access to data and facilities in the future could be a key tool to ensure accountability.

- **Where does the accountability lie:** In projects where data is collected by machines, and machines can even provide part of the data management plan where does that accountability sit? With the machine? The algorithm? The person who wrote the algorithm? These issues need to be discussed further.
- **Where does the raw data start:** A lot of data collection is now done by machines, and in some cases, even the "raw" data that comes off these machines isn't the actual raw data, because it has already gone through some initial pre-processing, e.g., to filter out "noisy data". This raises several issues, e.g., how raw should this raw data be? What noisy data should be kept in case important information is being lost? Further, we should be looking to harmonise these automated data collection processes as different machines and facilities could end up with vastly different looking data due to their different methods, which could prove extremely problematic if scientists were using multiple machines/facilities in their experiments.
- **How data is going to be collected and processed:** It is important for researchers to describe how they are going to collect and process their data. For example, if data is being collected via an audio recording, how secure is this? And what software are they going to be using to process this? Even if the data is collected securely, for example via an encrypted recording device in the same room as the person, this data could then be at risk once it is run through transcription software. This is why the entire lifecycle of the data needs to be thought through from the beginning.
- **Identify special category data:** Researchers need to really think about the special category data they are going to be collecting and describe it in detail.
- **Use FAIR principles:** Researchers should be helping the facilities they work in improve their metadata. Collected data should have appropriate metadata to describe it and make it understandable. Enough information should be given about the data and its collection to enable this process to be replicated, and the data should be reuseable.
- **Peer review and pride breeds accountability:** There is a competitiveness and pride associated with not wanting papers to be retracted which makes researchers accountable through that process. This makes it even more important to maintain.

Confidentiality

Confidentiality was picked because the levels of confidentiality required depend heavily on the data you are capturing. For example, medical records need a high level of confidentiality, but measuring rain via gauges in a field would have less strict requirements. The confidentiality levels need to appropriately match the data.

Considering this card also raised some further discussion points:

- **Be aware of human error:** Collecting data is a big responsibility, and the humans involved in data collection need to be wary of human error. Whether that is accidentally collecting the wrong data, or sharing confidential data inappropriately by accident.
- **Be aware of unforeseen consequences:** Sometimes the data you are collecting can inadvertently give away other things. For example sensors collecting data in a laboratory could inadvertently give away when the laboratory is occupied or empty; this is information that could be misused. Being aware of these types of edge cases would help to mitigate this. Further, applications can be hacked and misused if they rely on collecting data in a certain type of way. For example Google Maps uses location on mobile phones to build up a picture of the traffic in an area, but in February 2020 an artist used 99 phones to fake a traffic jam by wheeling them all over the same point at once. This was

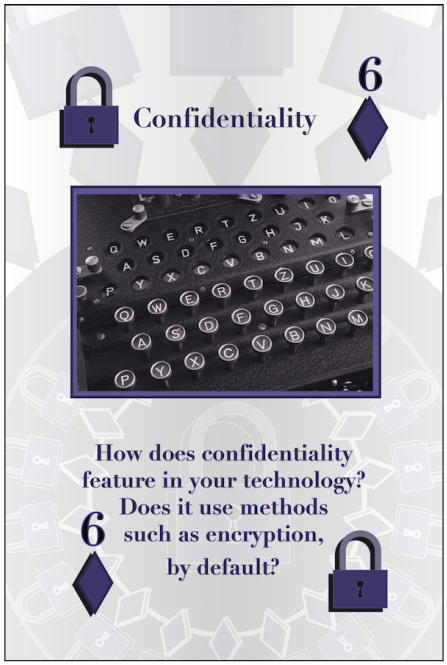

Figure 11: Confidentiality Moral-IT Card

done to demonstrate the flaws in the everyday systems that we take for granted.[62]

- **Users need to be careful too:** Users need to be reminded to be careful with the personal data they give away, and have a better understanding of the potential implications. There are some apps that collect location based data that users need to be careful with. For example, fitness apps such as strava log locations and times of your runs. They do provide security measures for users such as different levels of privacy settings and exclusion zones to mask the true location of your start and end points (typically a users house) but unless users are made aware of these options and understand why they are important, there are still risks here. This can also pose inadvertent risk to others, there have been cases of people running around air force bases as part of the military, which would inadvertently build up excellent maps of those locations.
- **Benefit outweighs concern:** There is a culture of giving away data for an easy life. It has become almost accepted that large companies such as Google, Amazon and Apple collect a plethora of data on you, but they provide services to make life easier which seems to outweigh this concern. For example, Amazon might collect a lot of data on me, but if I use their services I can get a parcel delivered by tomorrow morning.
- **We need more transparency:** We have progressed to a stage of unknown unknowns, we suspect that certain data is being captured. For example, our phones might be listening to our conversations, but we don't know for sure, or if it is happening, what the data is being used for. Companies need to be more transparent about this!
- **We are too overwhelmed to process:** Are there too many products out there for us to process? Meaning that we are happy to receive targeted ads if they are more relevant to us? Or are we happy for a business to take our details if they make switching a product easier? We need to delve deeper into these issues and understand the consequences of the choices that are made in these situations.

[62] https://www.wired.com/story/99-phones-fake-google-maps-traffic-jam/

5.4 Data Analysis

Data analysis is the part of the research lifecycle that pertains to using the data collected to discover new information and form conclusions. This often involves data cleaning, processing, transforming and modelling. The planned data analysis methods should be detailed in the data management plan.

The cards that were picked for this section were:

- Confidentiality
- Legibility & Comprehension

Confidentiality

Confidentiality was picked because there have been issues with researchers needing to conduct their interviews remotely due to COVID-19. Zoom is considered a very useful tool for this, as it provides automatic transcriptions. However it is unknown exactly where this data is stored, and the privacy laws surrounding this will differ depending on what country the user is in. Further, Zoom is an american company and therefore outside of GDPR and not covered by UK legislation.

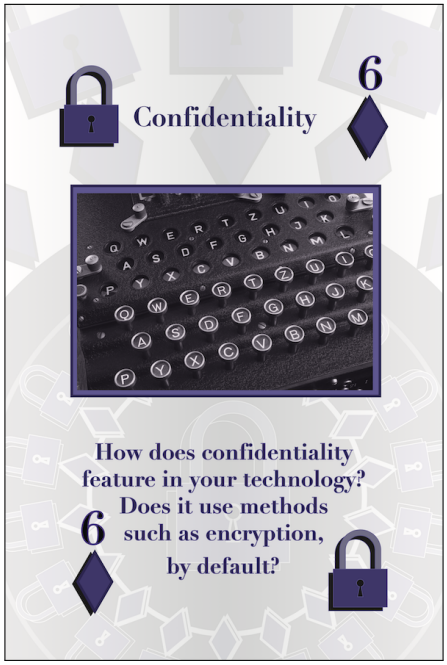

Figure 12: Confidentiality Moral-IT Card

Considering this card also raised some further discussion points:

- **Increase awareness of confidentiality:** The group agreed that confidentiality was a real issue with respect to data analysis, and that there was a clear need for further training in this area such that confidentiality requirements for data analysis could be taken into account before the data collection stage.
- **Researchers should talk to Data Managers more often:** It was very challenging to track whether advice given to researchers in this area (e.g., in how to manage confidentiality in their data) is actually taken on board. This applies to the different stages of the research data lifecycle. Other examples were given around getting involved with the data management planning stage, but then two years down the line the progress of the

project suggested that the advice given hadn't been taken on board and by that point it is too late. It would be useful to have more regular contact throughout a project to see if projects were actually following the advice they were given on their data management.

- **Data management plans aren't just for the application stage:** They need to stop being viewed as something that just needs to be created as part of an application, and start being viewed as an integral part of the project that should be adhered to even once you have reached the analysis stage.
- **Better case management:** There should be systems in place to track how research projects are progressing throughout, such that funding and publishing information can be tied in with them and their data management can be assessed more thoroughly throughout.

Legibility & Comprehension

Legibility & Comprehension was picked because there needs to be a better understanding of how data is actually produced in the first place, and how it has been analysed. It is too common to just see the end result of an analysis with no way to access the underlying data or attempt to reproduce the analysis methods used.

Figure 13: Legibility & Comprehension Moral-IT Card

Considering this card also raised some further discussion points:

- **Preserve processes and analysis:** There is a risk that by the final publication, there is a lack of information to understand how the data has got to the stage it has. Sometimes raw data is published, but other times it is the analysed data, which can be hard to understand without further information. Additionally, lots of information can be lost between paper notebooks and the formal write-up. We should be looking to capture and preserve the analysis methods to make our research more explainable and reproducible.
- **Encourage use of digital notebooks:** Electronic Lab Notebooks provide details about experiments and could be used to aid with understanding the processes and methods of analysis. These could be stored alongside the publication and data [23]. Although there are concerns surrounding risk of loss of intellectual property/ideas.
- **Further requirement for digital notebook tools:** There are lots of discipline specific

e-notebooks (especially across the sciences), but there is a need for a more generic version which could be used by other disciplines and shared across disciplines.

5.5 Data Sharing

Data sharing refers to the practice of making research data available to others. This is often a requirement from funders for research projects to ensure openness, transparency and reproducibility. How the data is going to be shared and in what form should be considered at the beginning of a research project and detailed in the data management plan.

There were three cards picked for this topic:

- Obfuscation
- Intellectual Property
- Accessibility

Obfuscation

Obfuscation was picked because there are often many questions with respect to anonymisation of data before sharing, such as how to anonymise data? And to what level?

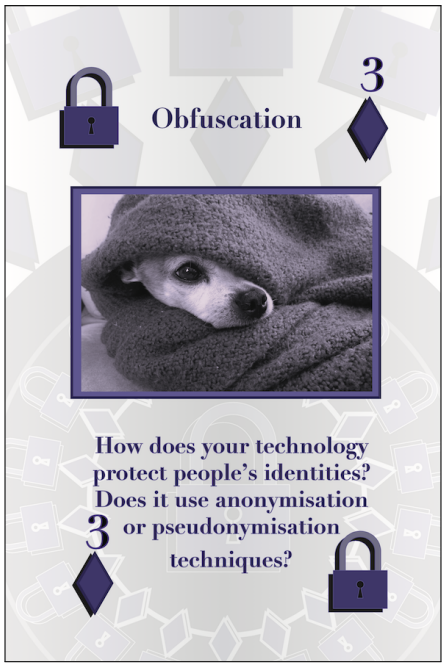

Figure 14: Obfuscation Moral-IT Card

Considering this card also raised some further discussion points:

- **There are never enough lawyers:** There aren't enough lawyers out there to support this work. People are trying to develop courses and guidance without the support of legal teams, when in reality this is a very complex matter.
- **There are issues around guidance for anonymisation:** Courses are being designed around this area, although there is a conflict between the requirement to provide general guidance and the need to serve researchers with particular backgrounds. Further, there is a requirement to ensure that researchers don't just end up with a checklist that means they don't critically engage with the whole project. Ideally researchers should be equipped with the tools to critically evaluate their own projects and anonymise where necessary.

There are some good resources in this area, for example the UK Anonymisation Network, but there is so much detail it can be hard to distil down into concise advice.

- **Tensions between disciplines:** Sometimes it can be most beneficial to refer researchers to experts working in their own domain as their advice may be more relevant than talking to someone from a completely different discipline.
- **Responsibilities of repositories:** Those in charge of the repositories also have a responsibility to ensure that data is redacted if it is personal data, and to ensure that the data submitted to them has the required level of anonymisation. However, this can then cause an issue of people patching fixes on at the last minute, which isn't ideal.

Intellectual Property

Intellectual Property was chosen because this can be a real issue with respect to data sharing, particularly if research projects involve commercial data/interests.

Figure 15: Intellectual Property Moral-IT Card

Considering this card also raised some further discussion points:

- **Start early:** There is often a tendency to only consider this at the end of a research project when data sharing comes to the forefront. This needs to be considered and planned for at the very beginning.
- **Clashes between industry and academia:** There can often be clashes between a desire to make research and data open access but still protecting the commercial aspects of research. These discussions need to be had at the beginning and IP statements/agreements put in place.
- **More awareness is needed:** Researchers often lack expertise in understanding what different licenses mean. There needs to be a greater awareness and understanding of these, and what might be "given away" by applying different licenses. It is worth consulting your libraries or departments that have expertise in copyright and licensing laws for advice.
- **Open Source doesn't mean give it away:** Data and software can both be open source but licenses can still be applied to them to dictate how and when they can be used, so it is possible to open things up without giving away all of your rights.

- **Consider both data and software:** You should consider the sort of licenses that your data needs as well as your software, as this is equally if not more important depending on the nature of the research.
- **Consider commercialisation:** It is worth considering from the beginning if you might want to commercialise your research, and making decisions accordingly due to that.
- **Tensions between openness and transparency / secrecy:** There is typically a standard three year embargo process for research projects, and then the data is required to be made open, often due to funding requirements. This gives researchers the space to do multiple experiments and publish their data. However, some researchers might want or need more time, and this needs to be given more thought, rather than it potentially just being released when the end date is reached. Especially as some researchers' careers can be dependent on generating these publications before others beat them to it. There should be some common justified reasons for not releasing things and for adapting these embargo times to meet individual project needs. Or there could be the potential for releasing "samples" of research data to make some of the project more accessible without giving away all of the details. This would also need to be considered further.

Accessibility

Accessibility was chosen because this can be a huge barrier to data sharing. Even if you have shared your data, if it isn't shared in a way that can be easily accessed and used by others then it somewhat defeats the point of sharing it.

Figure 16: Accessibility Moral-IT Card

Considering this card also raised some further discussion points:

- **Size matters:** The size of data can be a real issue in data sharing. If it is too big to be transferred over the internet then either you have to go to the data or the data has to come to you. This can have serious impacts, for example if it is far away, or how secure is it to travel with the data on a type of storage media compared to transferring it online? Is the data allowed to leave its current location? If not, how is this dealt with? There need to be more formal guidelines on how to handle this type of situation.

- **GDPR:** This legislation can impact the accessibility of data. At some facilities if you want to be granted access to a data catalog then you need to sign in and be a registered user. However, it is becoming increasingly common to be able to access some of these facilities anonymously because it reduces the data minimisation and reduces liability from a GDPR perspective because users' names aren't being stored etc. However this then means that you won't know who is accessing your data catalogs, so there is a trade off here.

6 Participants

Participants attended from a wide variety of backgrounds due to the online nature of the event. While the majority of attendees were from the UK, there were a number of registrations from other countries.

7 Conclusions

This event captured a wealth of information about the different ethical and legislative considerations that need to be taken into account when planning how to conduct your research and share your findings in a responsible way. A common recommendation from experts in these areas, and a repeated discussion output from the interactive activities across different parts of the research data lifecycle is that these issues need to be thought about on day one, rather than considered later down the line. As our technologies become ever more complex, and the wealth of information that can be gleaned from data (both intentionally and unintentionally) grows, we need to be even more mindful of planning our research projects ethically and in line with relevant legislation.

A key discussion point that arose in both the talks and interactive activity was the importance of data management plans, as often when things go wrong in research projects it is due to poor data management. These plans are the cornerstone of a successful research project, and they need to stop being viewed as something that is a tickbox exercise for an application, and start being viewed as an integral part of any research project that must be adhered to. Additionally, it was highlighted that we need to go further with our data management plans. These should be living documents that are updated throughout the project, and encouraging more contact with the experts who advise on data management plans such that they can ensure that the plans are being properly followed through. There is also now the concept of machine actionable data management plans to account for projects where a significant proportion of the data collected is generated by a machine. These options should be considered depending on the types of research projects you are planning.

All stages of the research data lifecycle need to be considered and planned for from the start. Researchers need to consider what data they actually need to collect, and look to only collect what is genuinely needed. How this data is going to be processed and anonymised should also be planned at the beginning, including factoring in the necessary resources and costs (whether that be just time or the cost for specialist software). Further, researchers should think about what they want to share, how it will be shared and who they're going to share it with. This helps ensure that when it is time to share and publish, the appropriate measures have been put in place.

Researchers also need to consider the edge cases of the data they could accidentally collect or even what unintentional information could be gleaned from it, even if it isn't being collected for that purpose. The discussions in this workshop gave several examples of how additional information could be calculated from data (e.g. working out where someone lives and regularly exercises by following their fitness data, or accidentally giving away secure locations by tracking

exercise around them) and these need to be properly thought through in advance. This is also a stark reminder that users of systems and users taking part in research studies also need to carefully consider what information they might accidentally be giving away. Thus there is a strong need to always inform participants of research studies very clearly about what data will be collected from them, and outlining exactly what they are agreeing to.

Finally, whilst this may sometimes seem like an insurmountable level of information for most researchers to get their heads around, there are some solid frameworks and principles that can and should be used to form a strong basis for how to manage research data responsibly. For example, GDPR has laid out a very solid base for responsible data management with their seven principles, and these should always be kept in mind and regularly returned to when considering processing any kind of personal data. Further, FAIR provides an additional set of principles for how to make your data findable, accessible, interoperable and reusable. There are also several sets of ethics guidelines that are in development such as the sets being developed by the UK, EU and United Nations. This report has provided links to a wide array of useful resources in this area in addition to summarising the key aspects of the talks and discussions that took place at this workshop, and we urge researchers planning prospective projects to read both these summaries and the recommended resources before embarking on creating their data management plans.

8 Related Events

Details of the other events in the Failed it to Nailed it data seminar series can be found here:
https://www.ai3sd.org/ai3sd-online-seminar-series/data-seminar-series-2020/
Each of these events will have video recordings and a report associated with it.

Details of other AI3SD events and events of interest can be found on the AI3SD website events page:
https://www.ai3sd.org/ai3sd-events/
https://www.ai3sd.org/events/events-of-interest/

References

[1] Beckles Z. AI3SD Video: Ethical data management – balancing individual privacy and public benefit;. AI3SD, PSDS & Patterns Failed it to Nailed it: Getting Data Sharing Right Seminar Series 2020. 2020. Available from: http://dx.doi.org/10.5258/SOTON/P0059.

[2] Narayanan A, Shmatikov V. Robust De-anonymization of Large Sparse Datasets. In: 2008 IEEE Symposium on Security and Privacy; 2008. p. 111–125. Available from: https://doi.org/10.1109/SP.2008.33.

[3] Rocher L, Hendrickx JM, de Montjoye YA. Estimating the success of re-identifications in incomplete datasets using generative models. Nature Communications. 2019 Jul;10(1):3069. Available from: https://www.nature.com/articles/s41467-019-10933-3.

[4] Voznik M. AI3SD Video: Data legislation, personal and non-personal data, ethical issues and protecting your IP rights;. AI3SD, PSDS & Patterns Failed it to Nailed it: Getting Data Sharing Right Seminar Series 2020. 2020. Available from: http://dx.doi.org/10.5258/SOTON/P0060.

[5] Kearns M, Roth A. The Ethical Algorithm: The Science of Socially Aware Algorithm Design. New York: Oxford University Press; 2020.

[6] Darbyshire T. AI3SD Video: Practical Ethics for Data Science and Algorithm Design;. AI3SD, PSDS & Patterns Failed it to Nailed it: Getting Data Sharing Right Seminar Series 2020. 2020. Available from: http://dx.doi.org/10.5258/SOTON/P0061.

[7] Munafò MR, Nosek BA, Bishop DV, Button KS, Chambers CD, Du Sert NP, et al. A manifesto for reproducible science. Nature human behaviour. 2017;1(1):1–9. Available from: https://doi.org/10.1038/s41562-016-0021.

[8] Martinez-Martin N. What are important ethical implications of using facial recognition technology in health care? AMA journal of ethics. 2019;21(2):E180. Available from: https://dx.doi.org/10.1001%2Famajethics.2019.180.

[9] Senior AW, Evans R, Jumper J, Kirkpatrick J, Sifre L, Green T, et al. Improved protein structure prediction using potentials from deep learning. Nature. 2020;577(7792):706–710. Available from: https://doi.org/10.1038/s41586-019-1923-7.

[10] Banerjee A, Prehoda E, Sidortsov R, Schelly C. Renewable, ethical? Assessing the energy justice potential of renewable electricity. AIMS Energy. 2017;5(5):768. Available from: http://dx.doi.org/10.3934/energy.2017.5.768.

[11] Hagendorff T. The ethics of AI ethics: An evaluation of guidelines. Minds and Machines. 2020;30(1):99–120. Available from: https://doi.org/10.1007/s11023-020-09517-8.

[12] Jobin A, Ienca M, Vayena E. The global landscape of AI ethics guidelines. Nature Machine Intelligence. 2019;1(9):389–399. Available from: https://doi.org/10.1038/s42256-019-0088-2.

[13] Diakopoulos N. Algorithmic Accountability. Digital Journalism. 2015;3(3):398–415. Available from: https://doi.org/10.1080/21670811.2014.976411.

[14] Coeckelbergh M. AI ethics. MIT Press; 2020.

[15] Bryson JJ. The Artificial Intelligence of the Ethics of Artificial Intelligence. The Oxford Handbook of Ethics of AI. 2020;p. 1.

[16] The Oxford Handbook of Ethics of AI: Online Supplement [Internet]; 2021. [cited 2021 Jan 25]. Available from: https://c4ejournal.net/the-oxford-handbook-of-ethics-of-ai-online-companion/.

[17] Bryson JJ. The Artificial Intelligence of the Ethics of Artificial Intelligence: An Introductory Overview for Law and Regulation. The Oxford Handbook of Ethics of Artificial Intelligence. 2019;.

[18] Cath C. Governing artificial intelligence: ethical, legal and technical opportunities and challenges. The Royal Society Publishing; 2018. Available from: https://doi.org/10.1098/rsta.2018.0080.

[19] Health TLD. Guiding better design and reporting of AI-intervention trials. Lancet Digit Health. 2020;2(10):e493. Available from: https://doi.org/10.1016/s2589-7500(20)30223-5.

[20] Raji ID. Handle with Care: Lessons for Data Science from Black Female Scholars. Patterns. 2020;1(8):100150. Available from: https://doi.org/10.1016/j.patter.2020.100150.

[21] Urquhart LD, Craigon PJ. The Moral-IT Deck: a tool for ethics by design. Journal of Responsible Innovation. 2021;p. 1–33. Available from: https://doi.org/10.1080/23299460.2021.1880112.

[22] Miksa T, Simms S, Mietchen D, Jones S. Ten principles for machine-actionable data management plans. PLoS computational biology. 2019;15(3):e1006750. Available from: https://doi.org/10.1371/journal.pcbi.1006750.

[23] Kanza S, Willoughby C, Gibbins N, Whitby R, Frey JG, Erjavec J, et al. Electronic lab notebooks: can they replace paper? Journal of cheminformatics. 2017;9(1):1–15. Available from: https://doi.org/10.1186/s13321-017-0221-3.