ORDNANCE SURVEY

PRACTICAL GeoAI ETHICS

Workshop output report 2022

Ordnance Survey

Copyright notice

Version history

Version	Date	Description
1.0	09/2022	Draft shared with workshop attendees for comment
1.1	12/2022	Final report

Contents

1. Foreword 4
2. Executive Summary 5
3. Introduction 7
4. Background 8
5. Workshop purpose 9
6. The workshop 10
7. Talks 12
7.1 Denise McKenzie - The Locus Charter: Ethics of Location 12
7.2 Jacqui Ayling – Moving from Principles to Practice 13
7.3 Jo Walsh & Mel Marochov – Model Cards: Our Journey as Developers to Incorporate Ethics into OS' GeoAI Workflows 14
7.4 Nigel Edmead & Vivek Sakhrani – Hyperlocal Mapping for Positive Change 15
7.5 Ruth Bowyer – Ethics and public perceptions of georesearch: learnings from a population-based cohort 16
7.6 Jeni Tennison – Public participation 17
8. Activities 18
8.1 Engaging with the Locus Charter principles 18
8.2 Practices that anticipate and mitigate harm 22
8.3 Public consultation, education, and local involvement 26
8.4 Active and continuous monitoring, assessment, feedback, explanation, and update of data and models 27
8.5 Frameworks for ethical conduct and project governance to be adhered to by all stakeholders 27
8.6 Data control and model control and governance 27
8.7 Data minimisation and anonymisation - only capture what is necessary 28
8.8 Summary 28
9. Discussion 30
10. Conclusion and next steps 32
11. Acknowledgements 33
12. Bibliography 34
13. Annexes 36

1. Foreword

By Jeremy Morley, Chief Geospatial Scientist, Ordnance Survey

The last few years have seen a welcome focus for Ordnance Survey on how our foundational data of the core geography of Great Britain should take account of a growing understanding in the field of location data ethics.

Through the Benchmark Initiative at Geovation, supported by Omidyar Network, we explored the ethical application of location data – starting a journey to think about what we at Ordnance Survey should learn and how we might contribute to the wider, global debate.

As the National Mapping Service of Great Britain, we provide authoritative geospatial data and services to HM Government, as well as a large range of customers both in the UK and internationally. We believe that ethically managed location data is critical to maintaining public, customer, and government trust, and demonstrating our organisational accountability.

Ordnance Survey signed the Locus Charter as a Supporter in September 2021 to make a long-term commitment to use location data responsibly and ethically, and encourage others to do the same. The Locus Charter was produced through collaboration, facilitated by The Benchmark Initiative and EthicalGeo, setting out a proposed set of common principles that can guide responsible practice when using location data.

The UK's Geospatial Commission has recently contributed to our understanding with a strategic interest in this area, undertaking an independent programme of public dialogue, and in June 2022, a policy paper describing how public confidence in location data may be built through ethical use.

We have started to learn from each other as to how the principles of the Locus Charter might apply in practice. The Association for Geographic Information (AGI) and the Royal Geographical Society (with IBG) convened key industry figures to continue the debate and build awareness of how we might start to adopt and/or adapt to take more transparent account of ethical practices in our work and in our data.

These discussions have been deeply inspiring, with individuals with a passion for location data, moving on our collective understandings as a geospatial industry, and as professionals and practitioners.

At Ordnance Survey, we wanted to explore further how to practise our principles, particularly in the field of GeoAI. Our research scientists worked collaboratively with customers, peers and key industry representatives to consider organisational risks, impacts, processes to help identify and avoid harms, and thinking about the practical tools we have available to us as organisations. We are publishing this report because we are truly grateful for those who contributed to our own understanding of ethical practices and we hope it will similarly be useful to others. As a community of organisations and professionals who live and breathe location data, we hope this report will be a stepping stone on our shared journey to improve the ethical use of location data and its applications.

2. Executive Summary

The ethical use of geospatial data and automated decision-making tools such as artificial intelligence (AI) have been a topic of discussion with increasing amplification over recent years. At Ordnance Survey (OS), we have been using machine learning to process geospatial (location) data for some time, and we are in the early stages of adapting practical tools to incorporate ethical principles into our work. The development of ethical principles, guidelines, and frameworks at the intersection between geospatial data and artificial intelligence (GeoAI) has laid important groundwork to inform and guide ethical practices. The next stage is to test, scale and implement practical tools throughout our GeoAI workflows.

To support our work and contribute to wider debate about location data ethics, Ordnance Survey brought together a number of key people to learn about existing tools that promote ethical practice in GeoAI and to identify the gaps where practical tools and processes are needed to help support the further development of ethical practices for those working at the intersection of GeoAI. We were – and remain – keen to share, collaborate and learn more about the current landscape of practically implementing ethics from others.

The workshop heard talks from data leaders and leading figures in the geospatial industry and academia. The workshop was an opportunity to engage with the Locus Charter principles of ethical use of location data, and to take stock and reflect on both the intersection of principle and practice and of Geo and AI. Practical examples already in use or being explored by organisations were highlighted which ought to anticipate and mitigate harms, for example by building trust by embedding data provenance, and ongoing work to encourage public participation in the control of their data.

Key learning points from the workshop:

- Developing **an 'ethics by design' mindset and culture requires leadership** which endorses ethical policies and practices. Clear lines of accountability within organisations and for customers to reach in to those organisations is essential. Structural changes need to develop in tandem with operational changes, and both support each other.
- **Principles need to be translated into practical processes and tailored** to the product, service, and organisation. That translation should be applied to actively guide ethical processes and design from project initiation (business/use case stage), through to development, deployment and monitoring. It is likely that implementing a range of practices within an organisation will be more effective than any single practice used in isolation.
- **Good ethical practice is highly nuanced**. Structures and processes will vary among organisations, but they would seek to identify ethical risks in some way, possibly across different functions and at different levels of an organisation. Practices might include: documentation of data and methods, organisational governance and internal connections, and engagement with those affected by the work. Stakeholder involvement and inclusive deliberation are essential if organisations are to identify and resolve complex ethical problems.
- More information and **detail about how to implement ethical practices is required**, allowing for variations between organisations, their disciplines and the embedding of ethical considerations into standard work pipelines. The development of tools and 'boilerplate' ethical methods to better define 'what good looks like' when taking these variances into account, could support the realisation of ethical GeoAI work.
- It may be that ethical practices will only truly standardise and coalesce with the introduction of regulation and legal frameworks. Ethical practices are sadly not the default for professionals,

practitioners and organisations working with location data today. Ethical practices have inevitable overheads and commitments to ongoing resource levels, and this can be a practical challenge. The commitment to continued dialogue within professions, cross-profession, and **the shared advocacy of principles can help build an environment which normalises 'ethics by design' as a core GeoAI practice**.

The workshop on practical GeoAI ethics was an effective spur to discussion about how we practise ethics within our organisations. As workshop participants at the time and post-event have said, there remains plenty more to work on. For instance, ethical GeoAI seemed to have more prevalent practice in the 'Geo' and less in 'AI', perhaps a reflection of the interests and backgrounds of the participants, and the greater 'knowability' of data in comparison to algorithms.

We would welcome contributions from AI practitioners into the discussion as it develops. There is a need to continue the discussion at the UK and international levels; to further collaborate on designing best practise to better define 'what good looks like'; and to continue to test how codes of professional and organisational principles can best support the application of ethics into business-as-usual processes, normalising 'ethics by design'.

3. Introduction

The ethical use of geospatial data and automated decision-making tools such as artificial intelligence (AI) have been a topic of discussion with increasing amplification over recent years. Increasing awareness within society of the use of data and data-processing tools to extract new knowledge has raised alarm among many. Even legal uses of data and AI could result in harm to individuals, communities, and non-human entities. For example, work that provides deeper information about locations of vulnerable or deprived communities can be essential for addressing the inequalities those communities experience but may also expose them to bad actors who wish to exploit those vulnerabilities. Thus, while these technologies can bring significant value to society, cultural, regulatory, and legal frameworks are needed to protect people and the world we live in.

Individuals and organisations in the data and technology sectors have a vital role to play in driving the conversation about ethics and risk, as well as facilitating good practice to avoid and mitigate potential harms. Much like good practice in health and safety, it is better to identify and avoid risks before they materialise. In doing so the benefits are two-fold; it is less likely that the actions of organisations or individuals cause harm to others, and the potential for reputational damage and subsequent negative impacts are reduced.

For this reason, several organisations are either designing or signing up to a set of ethical principles, such as the Locus Charter (Hawes & McKenzie, 2020), to which they pledge to align some or all aspects of their business. Committing to a set of principles is the first step along a journey that is currently poorly mapped. The next step is to implement structures, such as assigning ethically accountable roles within the organisation, and processes, like ethical risk assessment, to ensure that these principles are met. Further, such organisations need ways to provide evidence to internal and external stakeholders that they are accountable, responsible, and meeting these commitments.

As individuals within organisations, we find we have questions about how to practise our principles, such as:

- What structures and processes within organisations ensure that ethical risks are identified before they have negative impacts?
- How can these structures or processes be effectively implemented at all levels within an organisation?
- How should individuals adjust their practice to identify and avoid harms?

The aim of the Workshop on Practical GeoAI Ethics was to address such questions. This report brings together the materials used in the workshop, summarises the workshop outcomes, draws early conclusions as to practical tools available to organisations, and discusses the next steps in the field.

4. Background

The Locus Charter (Hawes & McKenzie, 2020) was developed in response to a recognised need for a set of common shared principles for the responsible use of geospatial and location data. The Benchmark Initiative, a collaboration between Geovation and the Omidyar Network, explored how the interests of the public can be protected whilst still encouraging the use of geospatial data for innovation. The resulting 10 founding principles (Figure 1: Locus Charter founding principles) were launched, in collaboration with EthicalGeo, on 24 March 2021. To date, at least 16 international organisations have signed up as Supporters of the Locus Charter (https://ethicalgeo.org/locus-charter/our-supporters/).

Figure 1: Locus Charter founding principles

As Ordnance Survey has signed up to the Locus Charter, as members of the research and innovation teams at Ordnance Survey, we are keen to understand what practical steps we could take within our roles to uphold the principles of the Charter. However, despite an extensive literature review on geospatial and AI ethics, we found ourselves uncertain what tools exist and what the best practice is. The idea of a workshop was born. At first, we assumed that we were behind the curve and that we could host a series of talks about different practices to help disseminate this knowledge. However, we quickly learned that many other organisations were at a similar moment – many were keen to adopt ethical frameworks but were uncertain what that might mean in practice.

We therefore designed a day that would encourage discussion, ideas sharing, and networking, with a view to drawing out themes and a range of solutions that participants could take back to their organisations for consideration.

The Workshop on Practical GeoAI Ethics took place on 16 June 2022 both in person, at Ordnance Survey headquarters in Southampton, and online over Microsoft Teams.

5. Workshop purpose

The workshop goals were to "*learn about existing tools that promote ethical practice in GeoAI*" and to "*identify the gaps where practical tools and processes are needed to help support the further development of ethical practices for those working at the intersection of GeoAI*".

We designed the workshop with the aim that participants would:

- gain a clear understanding of the Locus Charter
- contribute to a shared picture of existing practices of the ethical use of GeoAI
- collaboratively identify practical tools/methods for aligning their work with ethical principles
- receive a summary of the workshop (this report) after the event.

6. The workshop

The workshop was a mix of talks and activities hosted on Microsoft Teams, Slido, and Mural. Our ambition was to give all participants, both online and in-person, as similar experience as possible and use these tools to record interactions (but not identities). The agenda is shown in Figure 2.

10:30	Introduction | Isabel Sargent
	Slido poll: **Getting to know each other**
10:40	Talk: **The Locus Charter** | Denise McKenzie @PLACE
11:00	Refreshments
11:10	Mural activity: **Engaging with the Locus Charter principles** | Jacqui Ayling
11:45	Talk: **Moving from Principles to Practice** | Jacqui Ayling
12:00	Lunch
	Slido poll: **Where do we find ourselves?**
12:50	Talks: **Practical examples enacted in organisations** | Isabel Sargent
	Model Cards | Jo Walsh & Mel Marochov @Ordnance Survey
	Hyperlocal Mapping for Positive Change | Nigel Edmead @PLACE & Vivek Sakhrani @Atlas AI
	Ethics & Public Perceptions of Georesearch | Ruth Bowyer @The Alan Turing Institute & King's College London
	Public Participation | Jeni Tennison @Connected by Data
	Panel questions | Jeremy Morley
14:00	Refreshments
14:10	Mural activity: **Practices that anticipate and mitigate harm** | Isabel Sargent
15:30	Reflection: **Discussion with all participants** | Jeremy Morley
	Slido poll: **What will we take forward?**
16:00	Close

Figure 2: The workshop was a mixture of talks and activities

The first part of the day aimed to develop understanding of the Locus Charter, comprising a talk from Denise McKenzie and an activity in which participants were encouraged to consider and discuss the Locus Charter principles. The second part, with talks from Jacqui Ayling, Mel Marochov and Jo Walsh, Nigel Edmead and Vivek Sakhrani, Ruth Bowyer, and Jeni Tennison considered practical steps that organisations

could implement to operationalise ethics in their work. The final part of the day involved a quick-fire activity aimed at identifying possible ethical practices given a fictional geospatial mapping scenario, followed by open discussion.

Throughout the day, Slido was used to capture questions and launch 3 surveys to "take the temperature" of the participants of the workshop. Around 40 people attended with slightly more online than in-person. Breakout groups for the activities contained between 3 and 10 people; There were 3 or 4 groups online and 2 in person.

More details of each of these sessions are given in the following sections.

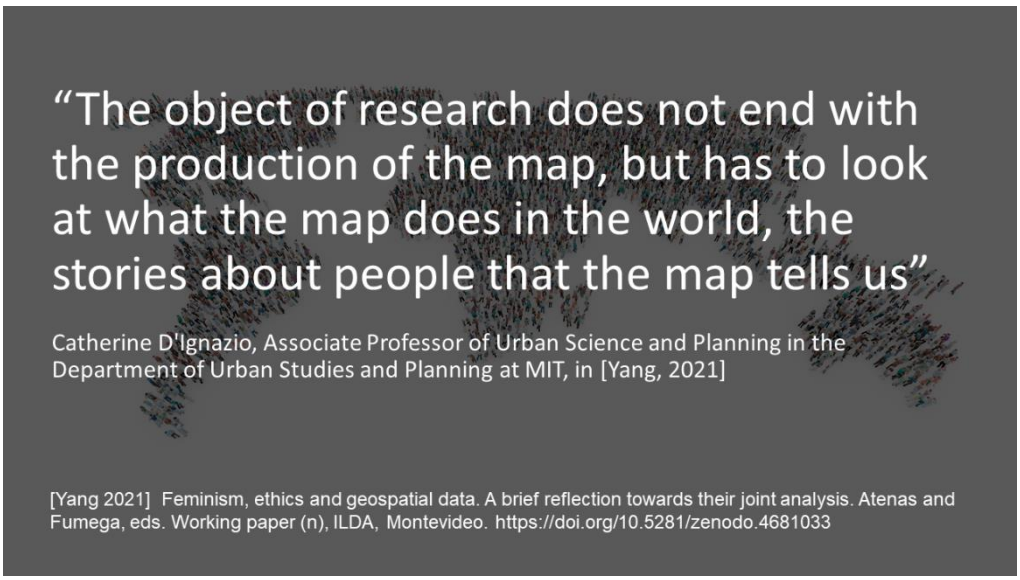

7. Talks

7.1 Denise McKenzie - The Locus Charter: Ethics of Location

Denise McKenzie from PLACE introduced the Locus Charter, which was developed with the Benchmark Initiative and EthicalGEO.

The Locus Charter is a proposed set of common international principles to support ethical and responsible practice when using location data. The Charter is written for individuals and organizations who use location data or have responsibility for activities that create, collect, analyse, and store location data. In her talk, Denise described how the Locus Charter came to be and how it encourages organisations to help protect individuals and the public interest.

Denise McKenzie is a strategic advisor, partnership builder, and presenter with over 20 years of experience with the global geospatial community. She works internationally to evangelize the benefits, value, and application of location data across government, the private sector, and academia and her experience covers a broad range of domains.

She co-directed the Benchmark Initiative operating through Ordnance Survey's Geovation accelerator, exploring the ethical use of location data, and it is through this role that she became co-author of the Locus Charter. In the broader geospatial community, she is the Chair-Emeritus of the board of directors for the Association for Geographic Information (AGI) in the UK and remains on council as the Lead for Ethics. Denise is also a member of the Global Advisory Board for the Location Based Marketing Association, and a steering committee member for Women in Geospatial+ leading their partnership program. Denise currently serves as Community and Ethics Partner at PLACE, see her bio to learn more.

The slides from this talk can be found in the Annex: A.

7.2 Jacqui Ayling – Moving from Principles to Practice

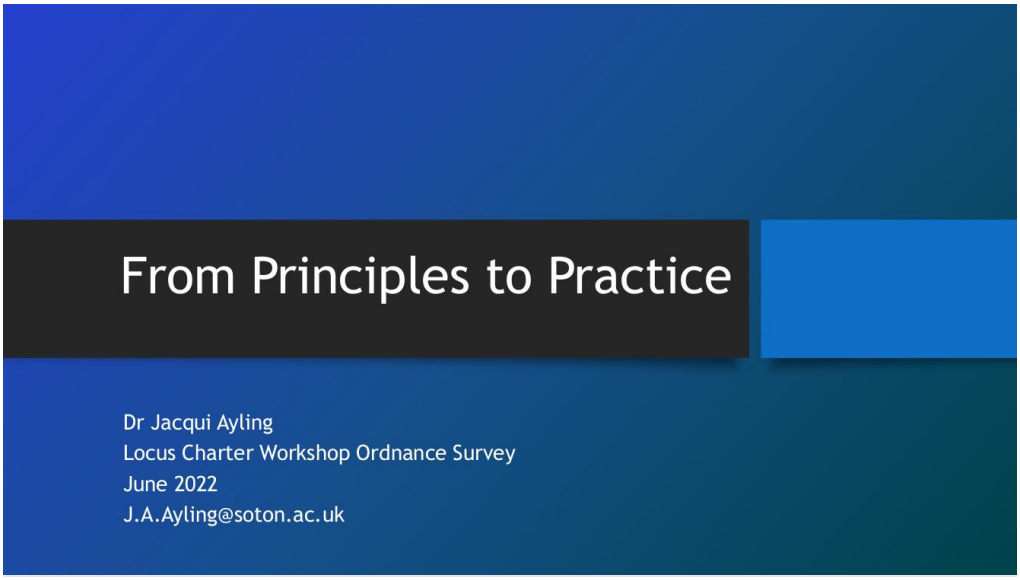

AI ethics has progressed considerably over the last few years. From identification of risks and harms, more and more organisations are signing up to ethical principles according to which they commit to operating. But there is still a great deal of work required to ensure that ethical principles are adhered to. Jacqui's contribution was focused on how organisations and individuals need to implement practices that ensure that they routinely identify, prevent, and mitigate ethical risks and harms.

This talk identified 8 steps to operationalising AI Ethics:

1. Commit to principles
2. Internal processes and controls
3. Designate accountability
4. Ethics committee or designated role
5. Technical tools – bias mitigation, model validation, documentation
6. Transparency
7. Whistle blowers – staff feedback
8. Stakeholder engagement

Dr Jacqui Ayling recently graduated with a PhD in AI and Data Ethics from the University of Southampton. She is passionate about translating ethical principles into practical application. She is continuing post-doc research in the governance implications of data provenance techniques for high-risk data flows, and leading an international research project on the emerging AI audit industry. She is currently disseminating the knowledge from her research by consulting on EU funded projects advising on ethics and governance of data institutions for energy communities. She is also teaching on professional training courses for the Southampton Data Science Academy. Jacqui has a background in environmental audit, bid writing, project development, and has global teaching and training experience in Higher Education.

The slides from this talk can be found in the Annex: B.

7.3 Jo Walsh & Mel Marochov – Model Cards: Our Journey as Developers to Incorporate Ethics into OS' GeoAI Workflows

In this talk, Jo and Mel walked through the journey of the VisionAI and Research teams at Ordnance Survey to incorporate ethical practice into their work at the developer level. They are beginning to use Model Cards to document how our machine learning models are made, how they can be used, and the ethical considerations that should go hand-in-hand with model development. The talk focused on their influences, ideas that they tested, the challenges, and their current work on an internal model hub – a central place for all model cards to live, which is designed to facilitate transparency, accountability and reduce duplication of effort.

Jo Walsh is a research software engineer and artist whose work on chatbots and the semantic web led to an interest in open geographic data. Jo is a former board member of the Open Knowledge Foundation and Open Source Geospatial Foundation, and a current board member of the Software Underground society of geoscience developers. At present Jo leads the VisionAI team at Ordnance Survey which offers in-house consultancy, research and support for deep learning data engineering.

Mel Marochov is a Graduate Data Scientist in the Rapid Prototyping Team (RPT) at Ordnance Survey. Before joining OS, she gained a Masters and published her work on using deep learning to classify satellite imagery of marine-terminating glaciers in Greenland. As a Graduate she has explored how OS can incorporate ethical practice into its GeoAI workflows and is currently enjoying the opportunity to learn everything she can from the RPT.

The slides from this talk can be found in the Annex: C.

7.4 Nigel Edmead & Vivek Sakhrani – Hyperlocal Mapping for Positive Change

This was a joint presentation by PLACE and Atlas AI. PLACE is a technology organization that collects and stewards high resolution optical imagery collected from the air and the ground using open-source, open component mapping drones and commercial mobile mapping systems in Sub Saharan Africa and small island states in the Caribbean. Atlas AI is a geospatial analytics company that uses satellite and aerial imagery, machine learning, and ground observations to generate high-resolution socio-economic data for data-sparse environments.

Nigel (PLACE) and Vivek (Atlas AI) presented on the organisational and technological approaches to geo-ethics. Nigel described the data PLACE collects as well as the governance, licensing and membership arrangements being established for the PLACE Trust, a legal trust to be established in the UK that will hold all PLACE data and licenses on behalf of PLACE supporters like Atlas AI. Vivek presented examples of Atlas AI's recent experiments on PLACE imagery to demonstrate the potential of ML techniques (and outcomes) using the high resolution imagery PLACE collects.

Nigel Edmead is a geo-focused learning development specialist with over 30 years' experience in the geo-spatial sector working in the UK, Africa and South East Asia. Nigel serves as Learning Partner at PLACE. Nigel also serves as Principal at enumanation, which provides learning advisory services to the geospatial sector and is an Associate of ConsultingWhere.

Vivek Sakhrani is Head of Analytics and Applied Data Science at Atlas AI and shapes Atlas AI's portfolio of development analytics services. He brings more than a decade of experience in systems planning, design, and investment advisory for development projects in energy, transport, water, ICT, and urban built environment. Vivek has a PhD in Systems Engineering and Master's in Technology and Policy both from MIT, where he also led research at the KACST-MIT Center for Complex Engineering Systems, the MIT Energy Initiative, and MIT Tata Center for Technology & Design. He is a member of the US National Academies' Transportation Research Board Urban Freight Committee, Assistant Editor for Engineering Project Organization Journal, and a board member for Ballroom Basix, a co-curricular dance and cultural program for K-12 students.

The slides from this talk can be found in the Annex: D.

7.5 Ruth Bowyer – Ethics and public perceptions of georesearch: learnings from a population-based cohort

Even where identifiable information is not being used, studies using geographic data should consider the ethics of their research on the population of the area their research is conducted on. Using a case study from the TwinsUK population-based cohort, Ruth presented reflections on how incorporating members of the public in research can improve it, and shared the participants perceptions of, and feeling towards, geotrace data.

Ruth Bowyer is an interdisciplinary researcher interested in systemic environmental influencers of human health across the life course. Her background spans biological, environmental, and ecological sciences, and she completed her PhD on environmental factors influencing the human gut microbiome in 2019. Her current interest lies in applying the 'OneHealth' framework, particularly in how climatic/environmental breakdown will influence health in ageing, multimorbidity and integrating geospatial questions into epidemiological studies. She currently works as a Research Associate at King's College London on COVID-19 related projects in the Department of Twin Research and Genetic Epidemiology, and at The Alan Turing Institute on climate change, heat & health. She has a keen interest in exploring different quantitative methodologies and finding innovative ways to communicate science, including in a grant-funded collaboration with an artist. In her spare time, she is an avid (if not particularly fast) runner and is most excited when a dog gets on the tube.

The slides from this talk can be found in the Annex: E.

7.6 Jeni Tennison – Public participation

Building trust with the public does not just involve putting into practice ethical principles, it means actively engaging them in the decisions you make about data. In this talk, Jeni described some of the ways in which organisations are bringing the public into the processes of data policy formation, operational decision making, and holding organisations to account.

Jeni Tennison is the founder of Connected by data, an initiative that aims to put community at the heart of data narratives, practices, and policies. She is co-chair of GPAI's Data Governance WG, a Shuttleworth Foundation Fellow and Associated Researcher at the Bennett Institute for Public Policy. She sits on the Boards of Creative Commons, the Global Partnership for Sustainable Development Data and the Information Law and Policy Centre. She was CEO of the Open Data Institute, where she worked for nine years. She loves Lego and board games and is the proud co-creator of the open data board game, Datopolis.

The slides from this talk can be found in the Annex: F.

8. Activities

8.1 Engaging with the Locus Charter principles

Following the talks, we wanted to engage the participants in an activity to relate the Locus Charter to their own work, to start to think about practice. In small groups, participants were asked to re-order the Locus Charter principles according to their collective view of each principle's importance, relative to the others. Each group had access to a small section of the Mural board illustrated in Figure 3. They could then move the coloured rectangle representing each principle from its initial position in the left-hand section to a new location on the right-hand section.

The purpose of this exercise was not to suggest that there is some 'correct' and fixed priority order to the principles. Instead, the intention was to provoke discussion about the different merits of the set of principles and their meaning and content.

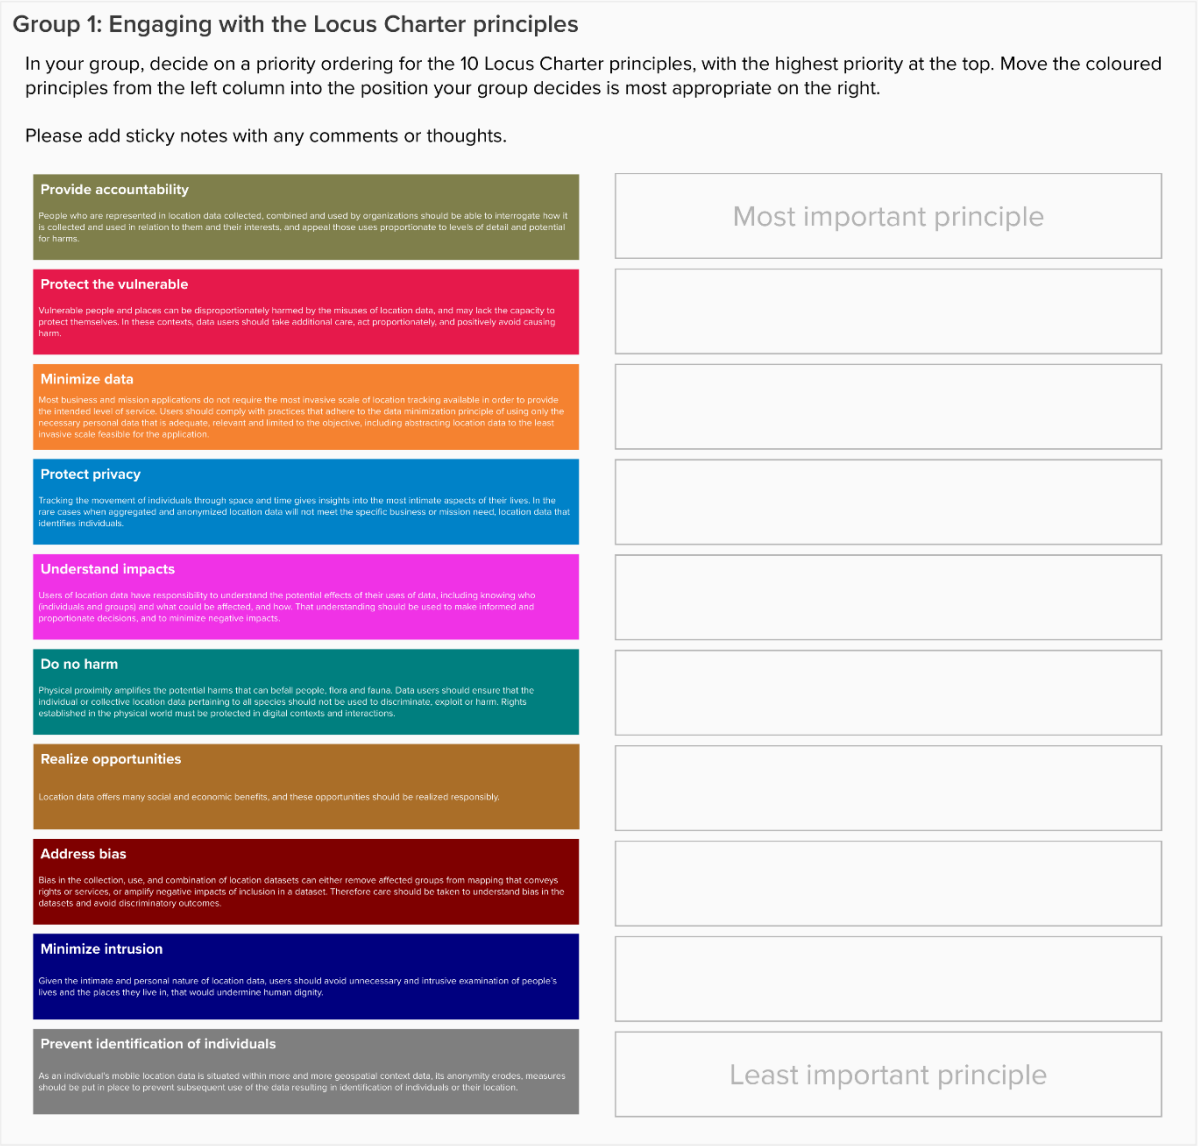

Figure 3: The initial layout of the activity board used by participants in the activity "Engaging with the Locus Charter principles"

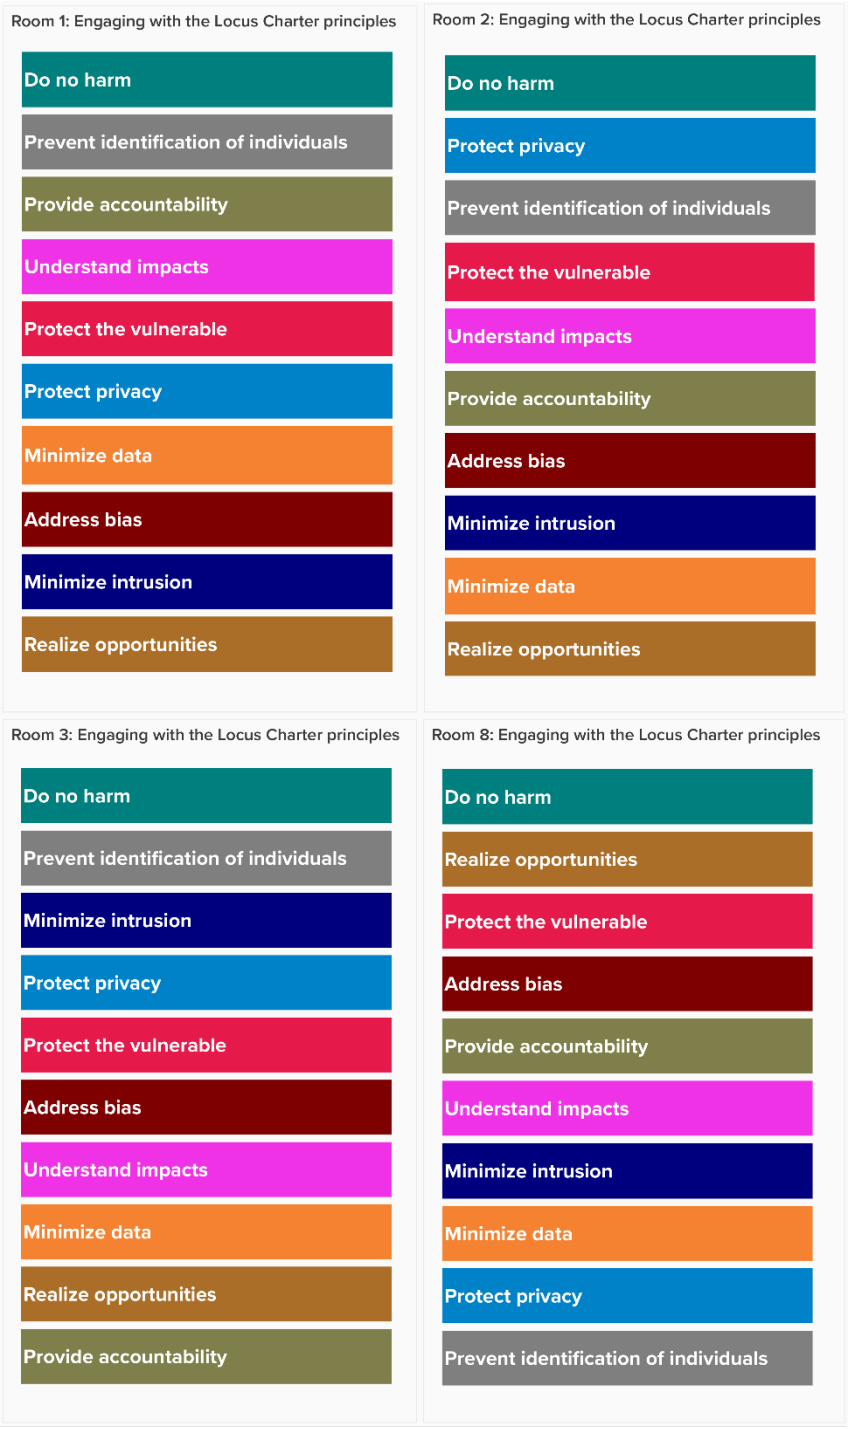

Figure 4: Four of the boards for the Locus Charter principle ordering exercise were completed

Figure 4 shows the four completed boards from this activity. There are quite a few similarities. All four chose to enter the "Do No Harm" as the most important principle with "Prevent Identification of Individuals" ranking highly on all but one of the boards. Interestingly, the principle "Realize Opportunities" ranked near the bottom for all but one of the boards (the same exception as earlier). The group who promoted this felt that it needed some emphasis because while 'doing nothing' might prevent harm, this

strategy might be harmful in net terms by not realising overall positive outcomes from the use of GeoAI technology.

Participants reflected that this exercise could be useful at the start and during a project, to consider what principles are most important during that work. With this in mind, it is less surprising that Realize Opportunities is ranked with lower importance, since one can assume that this is the purpose of any geospatial project.

It was recognised that there is some overlap between the principles (not least that "Do No Harm" essentially covers most of the rest in and of itself) so principles such as "Protect Privacy" and "Prevent Identification of Individuals" might not both need to be high in the ordering if one implied the other. This might explain the divergence in position of "Protect Privacy" across the four groupings.

It was also recognised that not all the principles will apply in any particular project – topographic mapping of building blocks from medium-resolution imagery using machine learning would, for example, be very unlikely to result in the identification of individuals.

8.2 Practices that anticipate and mitigate harm

Overview

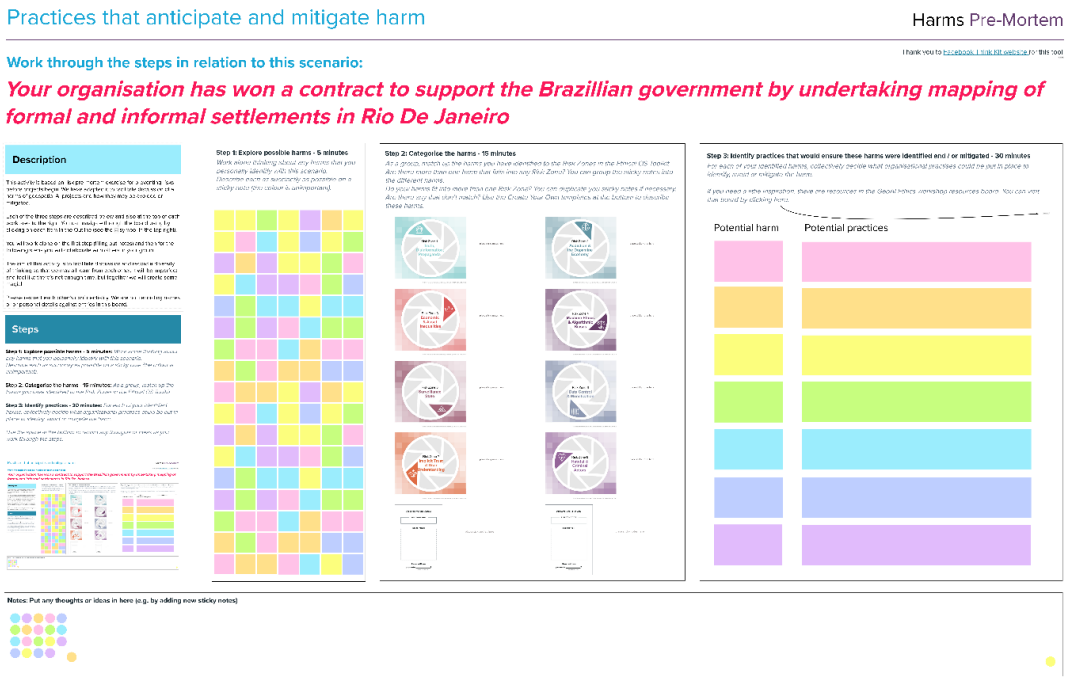

Figure 5: The initial layout of the activity board used by participants in the activity "Practices that anticipate and mitigate harm"

In this activity we attempted to rapidly design practices to identify and mitigate harms within a fictional case study "Your organisation has won a contract to support the Brazilian government by undertaking mapping of formal and informal settlements in Rio De Janeiro". This activity was designed to stimulate discussion, generate ideas, and to assimilate what we had heard earlier in the day.

Using a very simplified version of a project premortem (Eckert, 2015), participants first individually considered what harms could arise from the fictional mapping scenario. Then, in small groups, to assist grouping the identified harms together, participants matched them into Risk Zones using the Ethical OS Toolkit (Omidyar Network, n.d.). The final step in the activity was to share ideas about potential mitigation practices.

By assimilating the outputs from all the 5 breakout groups, 15 high-level harms arising from the fictional scenario became apparent:

- Breach of privacy
- Disclosure of vulnerable individuals / communities
- Lack of control / consent over use of personal data
- Data transfer / reselling (lack of control)
- Use of data to target crime, mistreatment, eviction, etc
- Use of data to damage landscape / environment
- Offense or insult caused by poor / no understanding of local culture

- Risk from mapping operation of mental harm to local people
- Increase in social inequality resulting from how the data are used
- Harm to reputation of organisation
- Harm to reputation of government
- Erroneous or biased data used in algorithms leading to erroneous outputs
- Erroneous outputs leading to bad conclusions / decisions
- Contribute to regime's propaganda by bestowing air of respectability
- Deliberate or accidental misinterpretation of data / misinformation

Figure 6 shows these potential harms and proposes a grouping into the entity that is harmed.

Participants also identified physical harms, which we would expect to be picked up by health and safety assessment in advance. This highlights how close physical risk assessment could be to ethical risk assessment.

These harms were distributed across seven of the eight Ethical OS Risk Zones, as well as possibly falling outside of these zones. Also, specific harms were found to arise within more than one Risk Zone.

Risk Zone 1: Truth, Disinformation, Propaganda:

- Breach of privacy
- Lack of control / consent over use of personal data
- Use of data to target crime, mistreatment, eviction, etc
- Deliberate or accidental misinterpretation of data / misinformation
- Contribute to regime's propaganda by bestowing air of respectability

Risk Zone 3: Economic & Asset Inequalities:

- Data transfer / reselling (lack of control)
- Use of data to damage landscape / environment
- Use of data to target crime, mistreatment, eviction, etc
- Increase in social inequality resulting from how the data are used
- Harm to reputation of government

Risk Zone 4: Machine Ethics & Algorithmic Biases:

- Erroneous or biased data used in algorithms leading to erroneous outputs
- Erroneous outputs leading to bad conclusions / decisions
- Increase in social inequality resulting from how the data are used

Risk Zone 5: Surveillance State:

- Data transfer / reselling (lack of control)
- Breach of privacy
- Disclosure of vulnerable individuals / communities
- Use of data to target crime, mistreatment, eviction, etc
- Harm to reputation of organisation

Risk Zone 6: Data Control & Monetization:

- Breach of privacy
- Lack of control / consent over use of personal data
- Data transfer / reselling (lack of control)

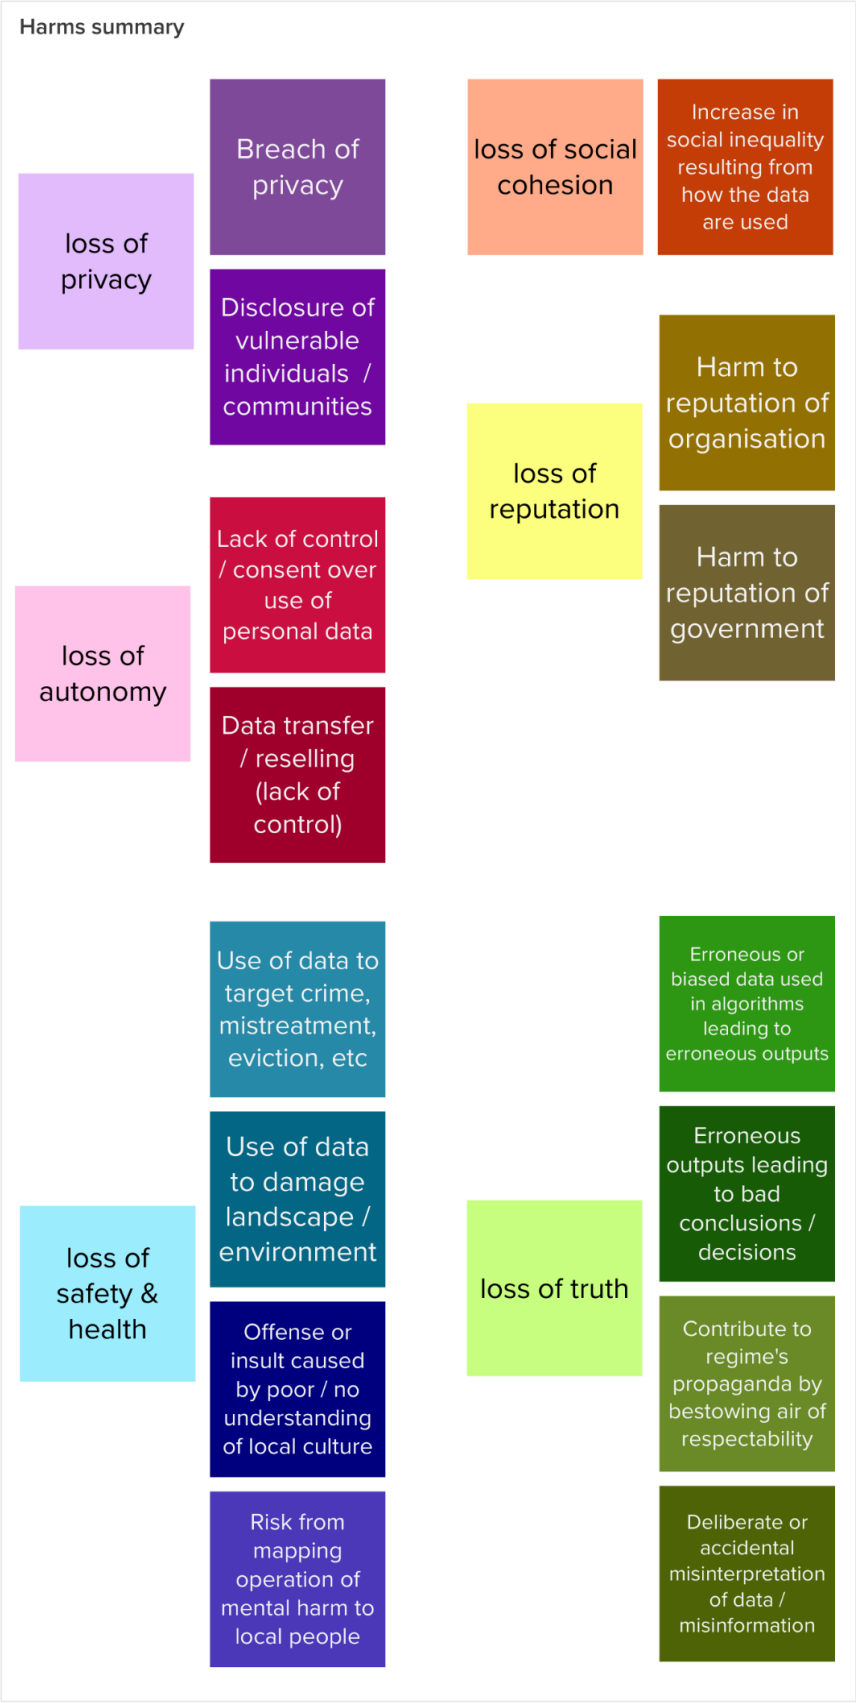

Figure 6: When activity outputs were grouped together, 15 harms were identified by participants in the workshop. These can be summarised into 6 different areas in terms of what is harmed.

Risk Zone 7: Implicit Trust & User Understanding:

- Breach of privacy
- Lack of control / consent over use of personal data
- Increase in social inequality resulting from how the data are used
- Risk from mapping operation of mental harm to local people

Risk Zone 8: Hateful & Criminal Actors:

- Lack of control / consent over use of personal data
- Use of data to target crime, mistreatment, eviction, etc
- Increase in social inequality resulting from how the data are used

Possibly outside of Risk Zones:

- Offense or insult caused by poor / no understanding of local culture
- Use of data to damage landscape / environment

Figure H1 to Figure H8 show all the harms identified, grouped into the different Ethical OS Risk Zones and Figure H8 shows harms that possibly fall outside of those Risk Zones.

The final step of this activity, to identify practices that would mitigate the identified harms, generated a great deal of thought and discussion. Despite the very short time available, 5 well-described areas of practice came out of the debate:

- public consultation, education, and local involvement
- active and continuous monitoring, assessment, feedback, explanation, and update of data and models
- frameworks for ethical conduct and project governance to be adhered to by all stakeholders
- data control and model control and governance
- data minimisation and anonymisation - only capture what is necessary

Each of these practice areas are described in more detail in the following sections. Figure 7 shows these 5 practice areas. In the appendices, Figure H9 shows all the suggested practices and how they summarise into these 5 areas of practice.

8.3 Public consultation, education, and local involvement

To address risks arising from lack of power of local populations and low awareness or understanding of the work being undertaken including potential loss of privacy, trust and harm to individuals, communities, and landscapes (which may not be apparent to the mapping agency), many groups proposed engaging the local (mapped) population in consultation activities, education, active participation, and local collaboration in the mapping process.

This would include providing information about how the data will be stored and shared, and having clear lines of data ownership before starting.

Such approaches were also identified to help counter risks arising from limited or biased data by introducing diversity and local knowledge to the work. In turn these would both help the identification and understanding of potential risks to local population and places, so that these may be avoided or mitigated, and could also help supply more diverse training data for any automatic process.

8.4 Active and continuous monitoring, assessment, feedback, explanation, and update of data and models

Harms that could arise from erroneous or biased data, and their resultant inference models, such as poor decision-making, loss of trust and reputation, and risks to physical, psychological, and environmental health were responded to with practices that ensure regular assessment of data and models.

This would include procedures for active monitoring and assessment / verification to identify bias or error, feedback processes to ensure that bias or error were remedied when discovered, and methods for data and model explainability / observability to ensure that decisions are transparent and taken within the scope of the domain of knowledge. Such practices would need to continue throughout the project and may involve outcomes of work with local populations to draw out risks that may not be apparent to other stakeholders.

8.5 Frameworks for ethical conduct and project governance to be adhered to by all stakeholders

Where risks arising from a political power imbalance, such as injurious actions inflicted (knowingly or otherwise) on the local population or environment, or disinformation that would disadvantage individuals or communities, many groups proposed drawing up clear missions and values, and ethical framework and transparent project governance framework to which all parties are signatories.

Such frameworks would outline clear accountabilities for all parties, with values including transparency, and provide the option to cease participation if other parties do not align with the stated values. It could extend to the collection, use and dissemination of the data and models created.

An additional motivation for such a framework was the identification that the organisation's reputation is at risk under these circumstances when deciding on which governments / other organizations you may want to work with gathering such data.

8.6 Data control and model control and governance

Setting out mechanisms of data and model control, before the work starts, was considered important to address the specific harms that may arise when data are repurposed, resold, or otherwise disseminated, and to ensure that all parties agree about a defined purpose for the data collection, creation, and use.

This would include a contractual licence with ethical framework describing accountability, ownership and authorisations, and the permitted uses of the data. It would also include transparent data lifecycle and

governance how the data will be stored and shared, and transparency around where the data are being collected.

Other control measures included creating working groups to control who specifically has access, to only share the results not the data processing, and to create linked data to reduce the opportunity to share. In a case where new uses of the data were identified, a further contract and framework would need to be drawn up.

8.7 Data minimisation and anonymisation - only capture what is necessary

Another practice that would reduce risk of harm to individuals, communities and environments that could arise from data creation and processing is to capture and produce strictly only what is required by the project and to anonymise any possibly personal information that remains.

This would require upfront and ongoing consideration of how necessary any data are before deciding to capture them. A data pipeline should be created that anonymised or generalised any personal or sensitive data that are required by the project.

8.8 Summary

This exercise, despite being undertaken in unrealistic circumstances, proved useful for identifying ways of mitigating the risks that could arise from the fictional mapping scenario. Exercises such as full-blown premortem (Eckert, 2015) could provide even more insight into possible harms and mitigation practices. The practices that were identified tended to group easily into areas of 'good practice' that should be incorporated into organisational culture and working habits, whether a particular project is being considered (or not).

All the practice areas incorporate a level of introspection which would mean that they would help maintain vigilance against new or worsening risks. Something that can seem so knotty about operating ethically is that there are so many unknowns. A developer may simply want to analyse data. A start-up may simply want to solve a customer's problem. The act of anticipating and mitigating harms that could arise further down the process may seem external to their roles' responsibilities. Therefore, ethical risk assessment needs to be part of business-as-usual, and considering the risk of harm resulting from one's actions must be part of daily practice.

Considering the 8 steps to operationalising AI Ethics, outlined in Jacqui Ayling's talk, the above practices fall into around half of the steps: "Internal processes and controls", "Technical tools", "transparency" and "Stakeholder engagement". If we assume that principles have been committed to, this leaves 3 further areas to address: "Designate accountability", "Ethics committee or designated role", and "Whistle blowers – staff feedback". These are structures that would be introduced by an organisation's leadership team.

Figure 7: The five practice types that were identified to mitigate harm during the second workshop activity

9. Discussion

The discussion covered several themes that had been raised during the day which could all be interpreted as practices, but perhaps not in the sense that we had designed the workshop.

An overriding theme was the importance of **building trust** by involving all stakeholders throughout the course of the project or activity. This is increasingly important because public trust with authorities is on the decline. This loss of trust can be attributed to poor engagement by actors such as big tech firms and (rightly) is resulting in more scrutiny of all uses of data and automatic decision-making.

Ethical practice cannot be ensured unless an **ethical culture** is created within an organisation. This requires accountability at board level, internal processes for staff to raise issues, and stakeholder engagement in organisational activities. Frequently, ethics is a niche conversation occurring in only some parts of organisations. Instead, ethical practice should be the default mode of operation and, moreover, outputs need also to be ethical 'by design'.

Of course, ethics are founded on the **values** of the organisation and the culture within which it operates. When capturing, processing, and interpreting data, it is rare to explicitly state the values underlying the activity because individuals are generally not cognisant of the values to which they are aligning. However, if social, epistemological, cultural, and other values are explicitly incorporated into GeoAI practise, for example the cost function for model optimisation, a breach of ethics may be less likely to occur. An important requirement here, going back to the first theme of building trust, is that to make values explicit within our work, requires considerable reflection and participation to ensure local, diverse, and inclusive input.

Some values are expressed within **legal frameworks**, but these rarely cover risk zones beyond physical harm and data control. Currently, much that is legal may be considered unethical, which leaves a large accountability gap if something goes wrong. As with data and privacy, it is likely that the great push for other aspects of ethics to be incorporated into corporate culture will only occur when they are regulated for.

It was noted that there had been far more discussion of location ethics and ethics in general than ethical practice in AI. AI presents very particular complexities because processes and models can lack transparency and happen at scales beyond the capacity of humans to monitor. It may therefore be necessary to consider if the principles of the Locus Charter are enough to cover AI practice.

The workshop considered ethical issues from many different perspectives: principles, stakeholders, potential harms, risk zones, and practice. The sense from participants was that, whilst there are indeed some people who are deeply conversant with ethics in technological fields, there are still many who are at the very start of the journey of understanding how to consider ethics within their field. One pertinent response to our question "what ethical practice would you like to have imitated in your organisation in the next year" was "Thinking about ethics in my work".

Considering our questions posed in the Introduction, the workshop talks showcased a range of practices that could highlight ethical risks and allow organisations to take mitigating action. These practices range from documentation techniques to be applied by those working with data and algorithms, to governance within and between organisations and practices that consult with, and engage those, affected by the organisation's activities. No doubt, best practise within an organisation would require a combination of practices co-ordinated across functions and at all levels within the organisation. Further, since all methods require creative thought to identify ethical risks, it will be important to ensure full engagement from all staff

and stakeholders to maximise the practices' effectiveness. However, currently ethical practice is a niche activity undertaken by an informed and self-motivated minority within organisations. If organisations are to align to their principles, practices must be embedded in the business-as -usual duties of all staff.

Finally, we will return to our workshop goals, which were to "*learn about existing tools that promote ethical practice in GeoAI*" and to "*identify the gaps where practical tools and processes are needed to help support the further development of ethical practices for those working at the intersection of GeoAI*". The workshop provided some good examples of existing tools particularly around stakeholder participation and data governance frameworks. The gaps are largely around the lack of corporate culture and structures and a legislative framework that would level the playing-field across all organisations. We look forward to many more discussions and hands-on activities that will fine-tune existing processes and structures, and produce new frameworks that will ensure GeoAI organisations can live the principles to which they subscribe.

10. Conclusion and next steps

The Workshop on Practical GeoAI Ethics was an effective spur to discussion about how we practise ethics within our organisations, and we believe gave many participants a greater breadth of understanding about how to ensure alignment to their chosen ethical principles. This breadth of understanding also highlighted that there are many different lenses through which ethical practice can be viewed. Jeni Tennison's blog post reviewing the event (Tennison, 2022), highlighted how there are nuances in ethical practice and that as much as we may try to "do no harm" harm will always be done somewhere because of an action, and it is a matter of choosing who benefits most. As Tom King, secretary of the Royal Statistical Society's Data Ethics Special Interest Group, said after the event "having the in-person component at the OS building was really good. A lot of ethics is normative and there is real value in gathering people and getting to try to reconcile different views and concerns, as several of the exercises did".

There remains plenty more to work on. For instance, the specific practices identified tended to be more in consideration of geospatial data, rather than AI algorithms, which is likely to reflect the interests and backgrounds of the participants, and the greater 'knowability' of data in comparison to algorithms. Also, the finer detail about implementing ethical practices still requires elaboration – and probably discovery.

This report, as well as describing the workshop event and the discussions that arose, marks a moment in time when many organisations are committing to ethical principles, and some are starting to think about what that means in terms of their practice. In the introduction, we said "Committing to a set of principles is the first step along a journey that is currently poorly mapped". We look forward to watching the journeys of different organisations as they move from pledging to act according to a set of principles to implementing ethics as part of their business-as-usual processes. We would like to see the discussion continue from the workshop, to see collaborate on designing best practise, so that others following from this moment can have a more clearly mapped path ahead of them.

11. Acknowledgements

This workshop could not have gone ahead without the support of Jacqui Ayling, who helped us develop content that encouraged discussion and interaction across the hybrid environment. Catherine Souch and Liz Fox-Tucker at RGS helped us consider ethics from a geospatial perspective, whilst Anne Steele, Jennifer Ding, Kirstie Whitaker and others at Alan Turing Institute helped us understand the different perspectives in the AI domain. Thank you to the presenters Denise McKenzie, Jacqui Ayling, Mel Marochov, Jo Walsh, Nigel Edmead, Vivek Sakhrani, Ruth Bowyer and Jeni Tennison who brought a wider breadth and depth of understanding to our discussions, in particular to Denise, Jacqui and Nigel who helped us realise through discussions in advance how useful this workshop would be. The day would have been so very chaotic without the calm assistance of Freya Cooper and Harry Baker who helped Mel Marochov in the online arena. Thank you also to Isabel Ulitzsch and the team at Geovation for helping us promote the event. Finally, thank you to the participants who engaged before the event, on the day and continue to keep up the momentum now.

12. Bibliography

Abdalla, M., & Abdalla, M. (2021). The Grey Hoodie Project: Big Tobacco, Big Tech, and the threat on academic integrity. *arXiv:2009.13676*. Retrieved from https://arxiv.org/abs/2009.13676

Anand, S., & Bailey, K. (2022, May 25). *How digital feudalism hurts farmers*. Retrieved July 5, 2022, from How digital feudalism hurts farmers: https://datavaluesdigest.substack.com/p/how-digital-feudalism-hurts-farmers

Ayling, J., & Chapman, A. (2020, September 12). Putting AI ethics to work: are the tools ft for purpose? *AI and Ethics*. doi:https://doi.org/10.1007/s43681-021-00084-x

Bird, E., Fox-Skelly, J., Jenner, N., Larbey, R., Weitkamp, E., & Winfield, A. (2020, March 1). The ethics of artificial intelligence. *The ethics of artificial intelligence*. Retrieved from https://www.europarl.europa.eu/RegData/etudes/STUD/2020/634452/EPRS_STU(2020)634452_EN.pdf

Eckert, T. (2015). The pre-mortem: An alternative method of predicting failure. *2015 IEEE Symposium on Product Compliance Engineering (ISPCE)*, (pp. 1-4). doi:10.1109/ISPCE.2015.7138700

Ethically Aligned Design. (2016). *Ethically Aligned Design*. Retrieved from https://ethicsinaction.ieee.org/

European Union. (2021). Proposal for a Regulation laying down harmonised rules on artificial intelligence. *Proposal for a Regulation laying down harmonised rules on artificial intelligence*. Retrieved from https://digital-strategy.ec.europa.eu/en/library/proposal-regulation-laying-down-harmonised-rules-artificial-intelligence

Gebru, T., Morgenstern, J., Vecchione, B., Vaughan, J. W., Wallach, H., au2, H. D., & Crawford, K. (2020). Datasheets for Datasets. Retrieved from https://arxiv.org/abs/1803.09010

Hawes, B., & McKenzie, D. (2020). *Locus Charter*. Retrieved February 4, 2021, from Locus Charter: https://ethicalgeo.org/wp-content/uploads/2021/03/Locus_Charter_March21.pdf

IEEE Standards Association. (2019). Ethically Aligned Design. *Ethically Aligned Design, First*. Retrieved from https://standards.ieee.org/content/ieee-standards/en/industry-connections/ec/ autonomous-systems.html

IEEE Standards Association. (2020, February 4). Ethically Aligned Design for Business. *Ethically Aligned Design for Business*. Retrieved from https://standards.ieee.org/content/dam/ieee-standards/standards/web/documents/other/ead/ead-for-business.pdf

Kelley, S. (2022, February 10). Employee Perceptions of the Effective Adoption of AI Principles. *Journal of Business Ethics*. Retrieved from https://link.springer.com/article/10.1007/s10551-022-05051-y

Leslie, D. (2019). *Understanding artificial intelligence ethics and safety: A guide for the responsible design and implementation of AI systems in the public sector.* Tech. rep., The Alan Turing Institute. doi:https://doi.org/10.5281/zenodo.3240529

Mitchell, M., Wu, S., Zaldivar, A., Barnes, P., Vasserman, L., Hutchinson, B., . . . Gebru, T. (2019, January). Model Cards for Model Reporting. *Proceedings of the Conference on Fairness, Accountability, and Transparency*. doi:10.1145/3287560.3287596

Nikuze, A., Flacke, J., Sliuzas, R., & van Maarseveen, M. (2022). Urban induced-displacement of informal settlement dwellers: A comparison of affected households' and planning officials' preferences for resettlement site attributes in Kigali, Rwanda. *Habitat International, 119*, 102489. doi:https://doi.org/10.1016/j.habitatint.2021.102489

Omidyar Network. (n.d.). *Ethical OS Toolkit*. Retrieved from Ethical OS Toolkit: https://ethicalos.org/

Open Data Institute. (2019, July 3). *ODI Data Ethics Canvas*. Retrieved February 4, 2021, from ODI Data Ethics Canvas: https://theodi.org/article/data-ethics-canvas

Strümke, I., Slavkovik, M., & Madai, V. I. (2021). The social dilemma in artificial intelligence development and why we have to solve it. *AI Ethics*. doi:https://doi.org/10.1007/s43681-021-00120-w

Syed, M. (2019). *Rebel Ideas.*

Tennison, J. (2022, June 16). *Workshop on Practical GeoAI Ethics.* Retrieved from Connected by Data: https://connectedbydata.org/events/2022-06-16-geoai-ethics

UK Statistics Authority. (n.d.). *Ethical considerations in the use of geospatial data for research and statistics*. Retrieved from Ethical considerations in the use of geospatial data for research and statistics: https://uksa.statisticsauthority.gov.uk/publication/ethical-considerations-in-the-use-of-geospatial-data-for-research-and-statistics/

Yang, S. (2021). *Feminism, ethics and geospatial data. A brief reflection towards their joint analysis.* Tech. rep., ILDA, Montevideo. doi:https://doi.org/10.5281/zenodo.4681033

13. Annexes

A. Slides presented by Denise McKenzie - The Locus Charter: Ethics of Location 37

B. Slides presented by Jacqui Ayling – Moving from Principles to Practice 45

C. Slides presented by Jo Walsh & Mel Marochov – Model Cards: Our Journey as Developers to Incorporate Ethics into OS' GeoAI Workflows 49

D. Slides presented by Nigel Edmead & Vivek Sakhrani – Hyperlocal Mapping for Positive Change 54

E. Slides presented by Ruth Bowyer – Ethics and public perceptions of georesearch: learnings from a population-based cohort 57

F. Slides presented by Jeni Tennison – Public participation 65

G. Slido surveys 71

H. Practices that identify and mitigate harms outputs 75

A. Slides presented by Denise McKenzie - The Locus Charter: Ethics of Location

PLACE

PLACE is a non-profit mapping organization founded on the belief that high quality mapping is essential data infrastructure for the digital age

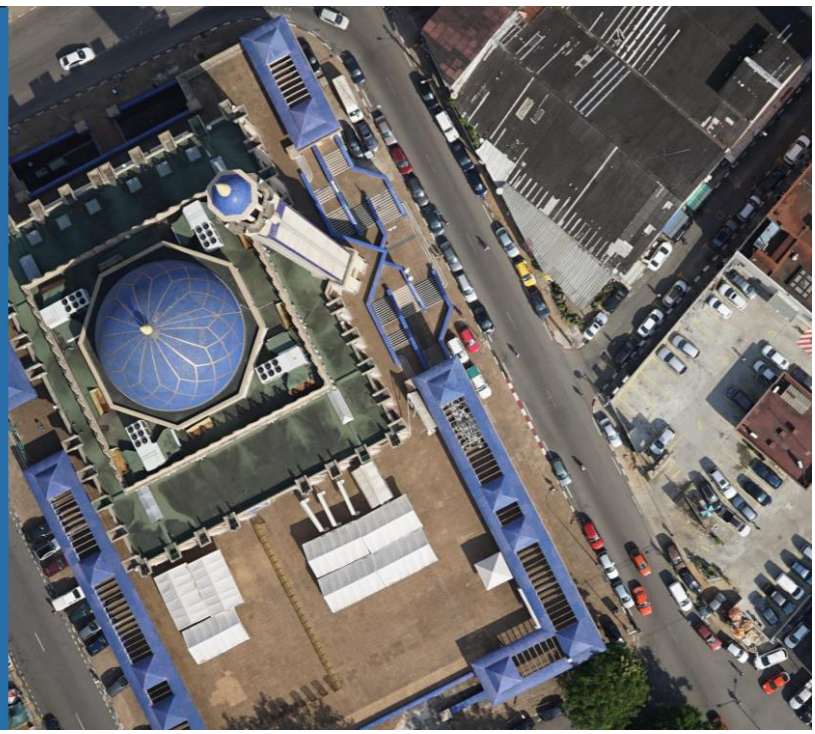

Our **MISSION** is to map the urban world in ultra-high resolution and make the images open, reliable and accessible and place them for perpetuity in a UK based legal trust

PLACE

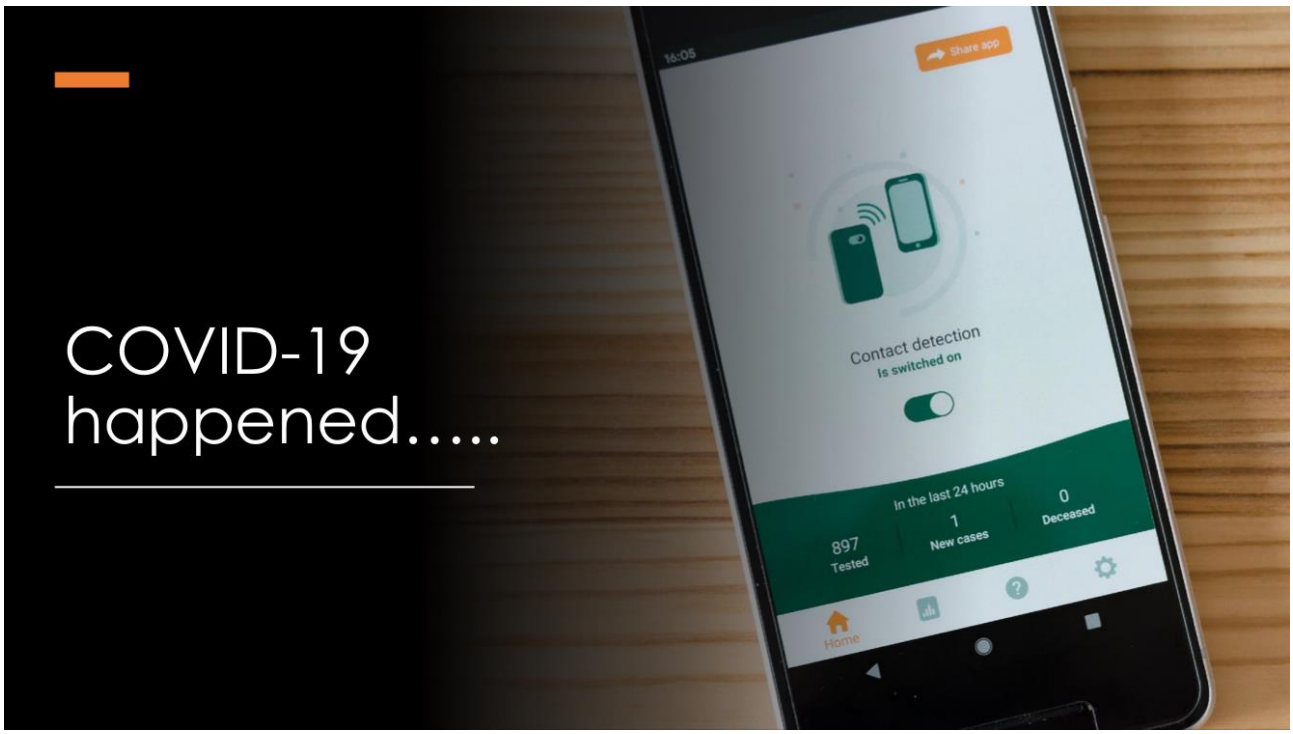
COVID-19
happened.....
16:05
Share app
Contact detection
Is switched on
In the last 24 hours
897
Tested
1
New cases
0
Deceased
Home

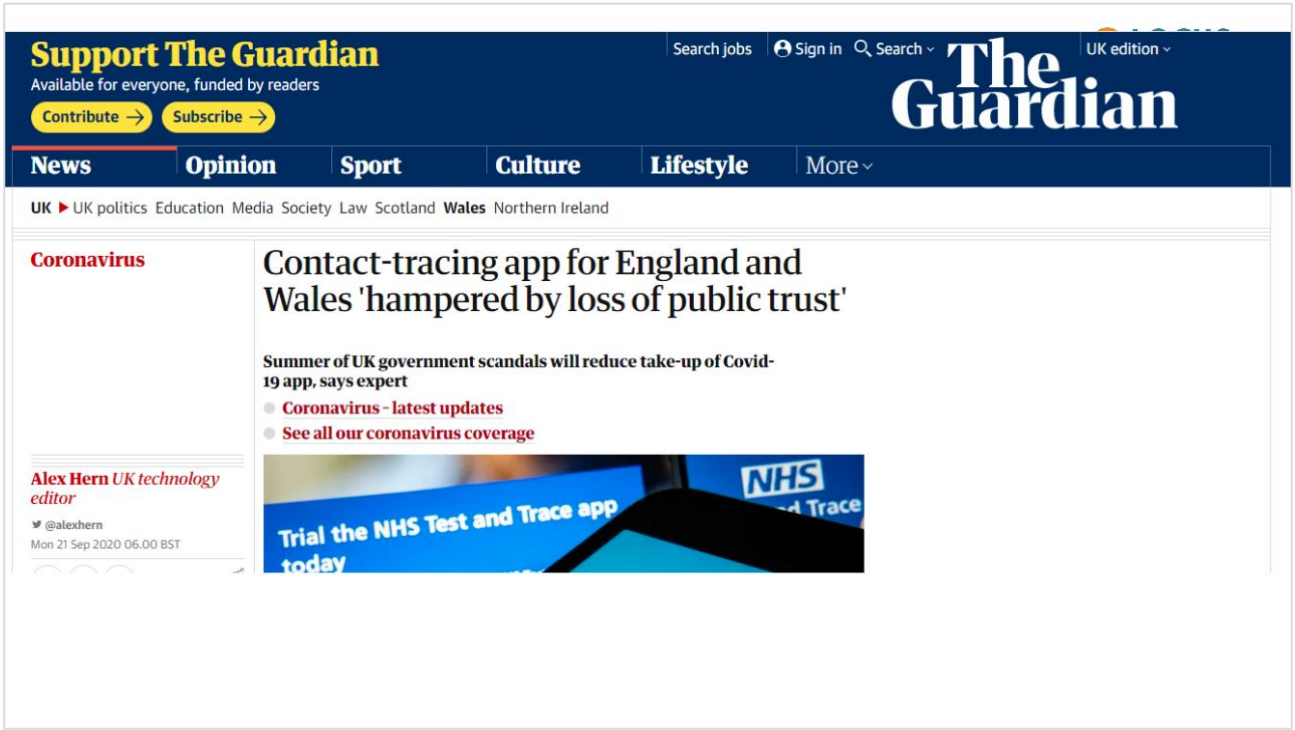
Support The Guardian
Available for everyone, funded by readers
Contribute
Subscribe
Search jobs
Sign in
Search
The Guardian
UK edition
News
Opinion
Sport
Culture
Lifestyle
More
UK
UK politics
Education
Media
Society
Law
Scotland
Wales
Northern Ireland
Coronavirus
Contact-tracing app for England and Wales 'hampered by loss of public trust'
Summer of UK government scandals will reduce take-up of Covid-19 app, says expert
Coronavirus - latest updates
See all our coronavirus coverage
Alex Hern UK technology editor
@alexhern
Mon 21 Sep 2020 06.00 BST
Trial the NHS Test and Trace app today
NHS

LOCUS CHARTER

"Datafication" expressing and managing the world with data

Bias - Privacy - Market power and data colonialism
Complexity - Transparency - Trust

Benchmark Initiative

- Supported by PLACE (part of the Omidyar Network) and Ordnance Survey
- Complementary program in USA EthicalGeo

BENCHMARK
Raising Standards for Location Integrity

GEOVATION

PLACEFUND

ON
OMIDYAR NETWORK

EthicalGEO

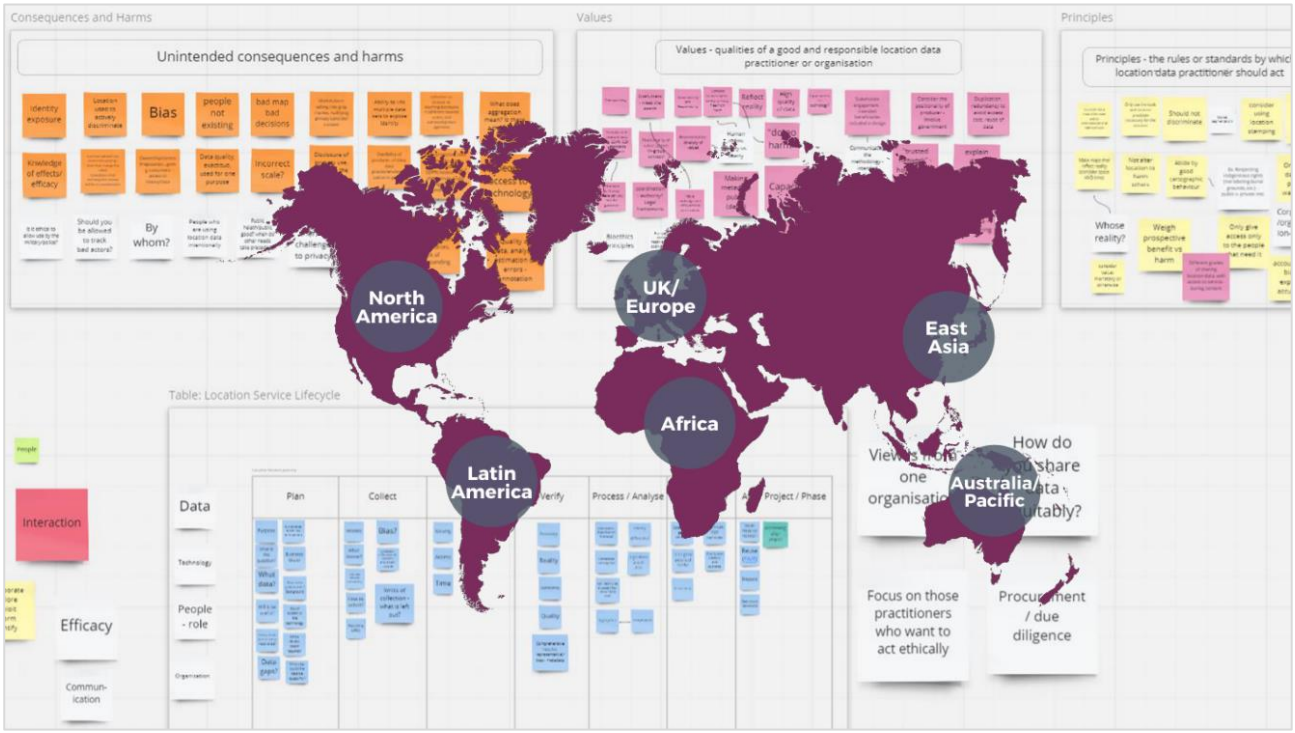

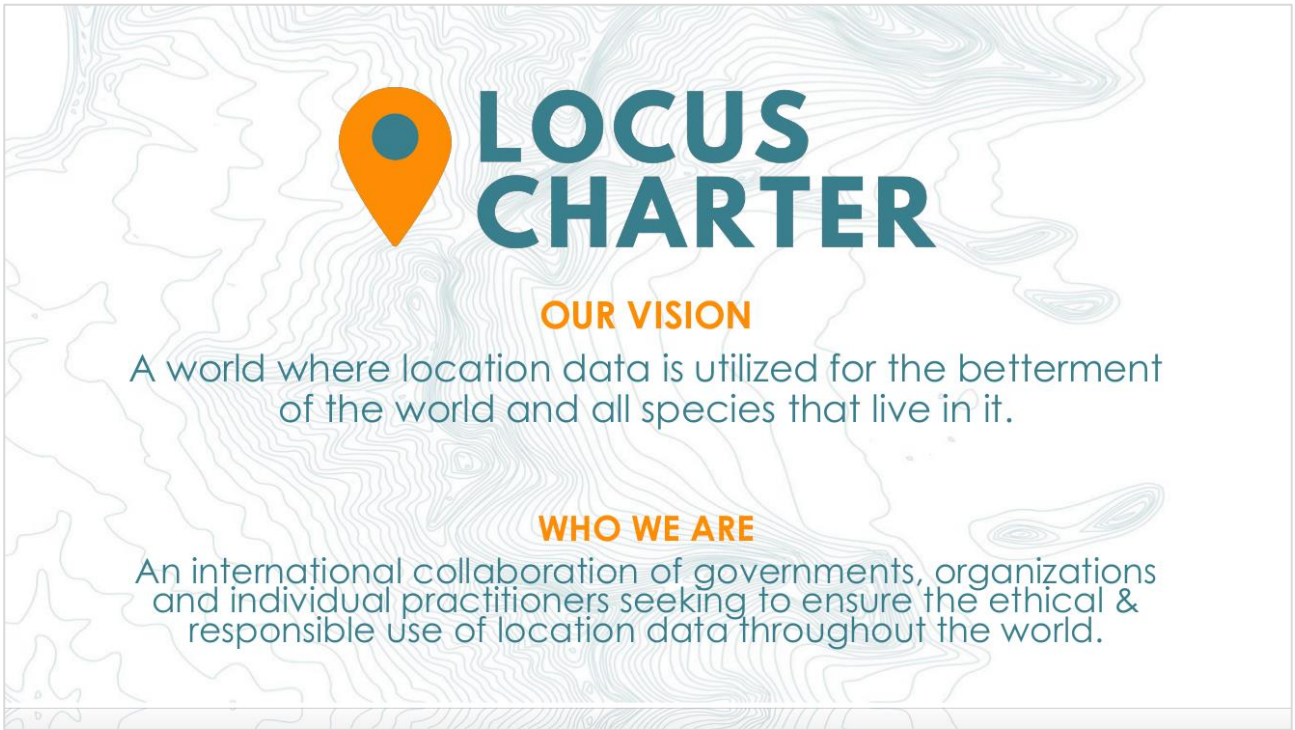

What can you do?

- Use the Locus Charter to structure questions about practice and impacts
- Use the Charter as a basis for roundtables, training and organisational development
- Join the Community and the international conversation
- https://ethicalgeo.org/locus-charter/
- #locuscharter
- Share related initiatives in ethical use of local data

Related initiatives

Activities	Document Type	Audience
Locus Charter (Benchmark Initiative & EthicalGEO)	Strategic Global Principles	Organisations (private & public) primarily, but can also be endorsed by individuals
ODI Data Ethics Canvas & Data Ethics Maturity Project GEO – Data Working Group Ethics best practice Geonovum – Ethical Framework OGC – GeoEthics adhoc *(proposed working group)* OECD – Geospatial Lab Ethics Workstream	Frameworks / Best Practices	Organisations
W3C SDWWG – Responsible Use Guide Godan – Ethical Code Toolkit SDSN TReNDS – Contracts for Data Collaboration Omidyar Network - Ethics Explorer DevGRG – Development Research Ethical Guidelines Geospatial Commission (UK) – Location Data Ethics Report	Guides / Guidelines / Templates	Practitioners implementing on a daily basis
URISA / GISCI (USA) SSSI (Australia & NZ) RICS (UK) ASPRS (USA)	Codes of Ethics	Individuals, Professionals

Be part of the community

Read the charter & join the community
https://ethicalgeo.org/locus-charter/

Contact the team info@ethicalgeo.org

Follow us on twitter @locuscharter

LOCUS CHARTER

B. Slides presented by Jacqui Ayling – Moving from Principles to Practice

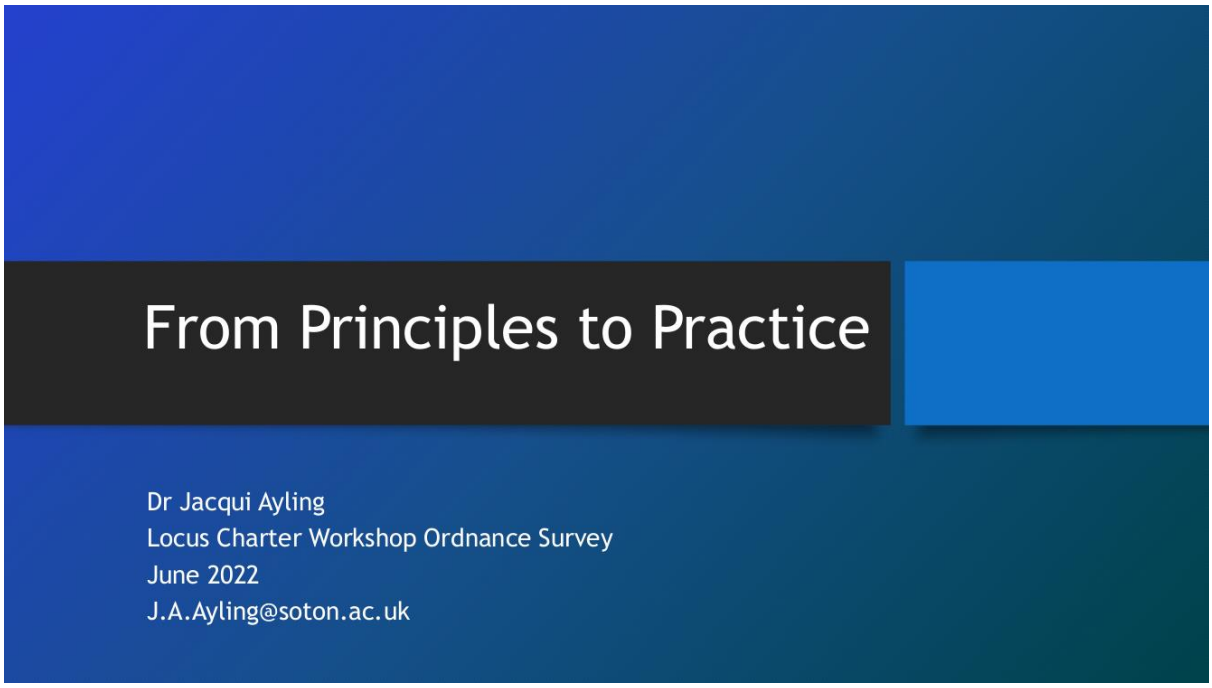

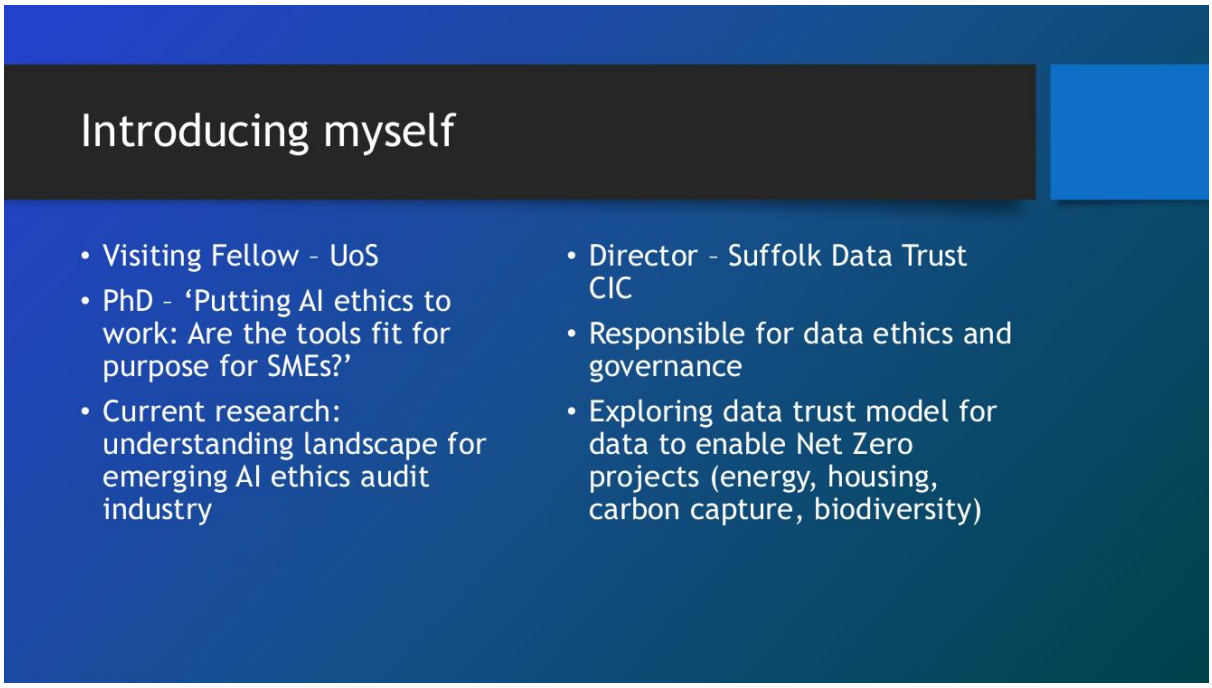

Who do we need to consider?

- Humans as:
- **employees**
- **citizens or subject to government**
- **clients, prospects or users**
- Also… non-humans? The environment?

Steps to operationalise AI ethics

1. Commit to principles	2. Internal processes and controls	3. Designate accountability	4. Ethics committee or designated role
5. Technical tools – bias mitigation, model validation, documentation	6. Transparency	7. Whistle blowers – staff feedback	8. Stakeholder engagement

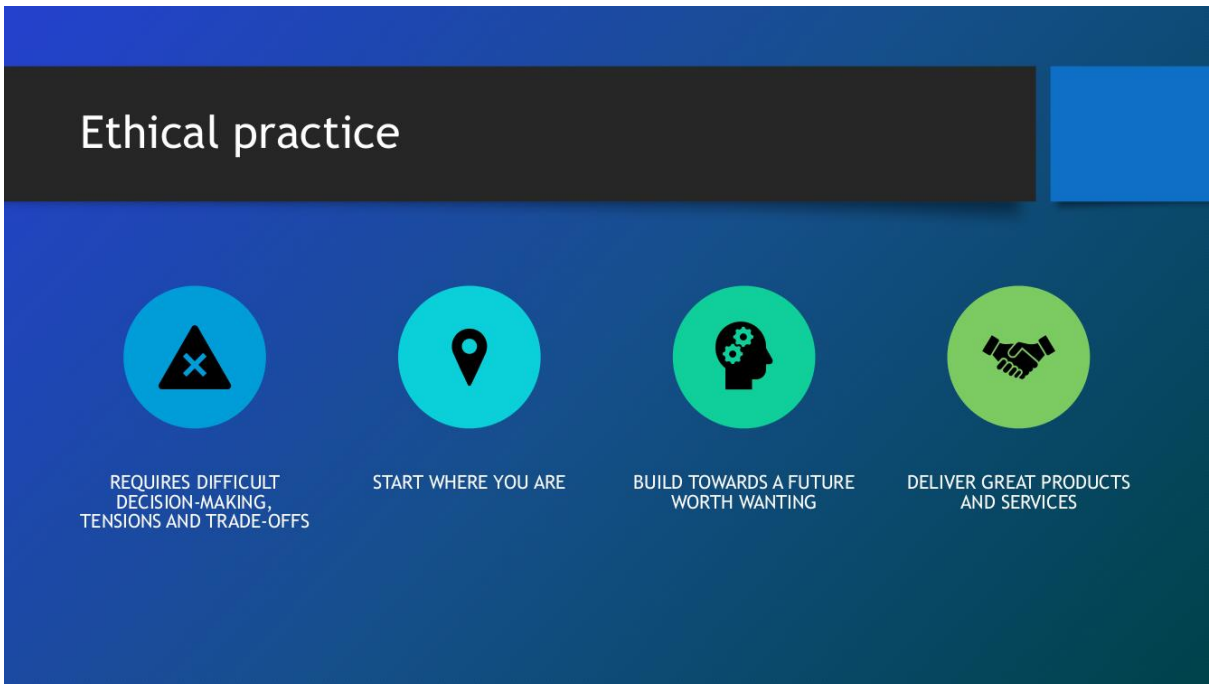

Further resources

- Ayling, J. and Chapman, A. (2021) 'Putting AI ethics to work: are the tools fit for purpose?', AI and Ethics. doi:10.1007/s43681-021-00084-x
- Doteveryone (2019) 'Consequence Scanning – doteveryone'. Available at: https://doteveryone.org.uk/project/consequence-scanning/
- Energy Systems Catapult (2021) 'Energy Data Management Canvas', Energy Systems Catapult. Available at: https://es.catapult.org.uk/guide/energy-data-management-canvas/
- Institute for the Future and Omidyar Network (2018) 'Ethical OS'. Available at: https://ethicalos.org/

C. Slides presented by Jo Walsh & Mel Marochov – Model Cards: Our Journey as Developers to Incorporate Ethics into OS' GeoAI Workflows

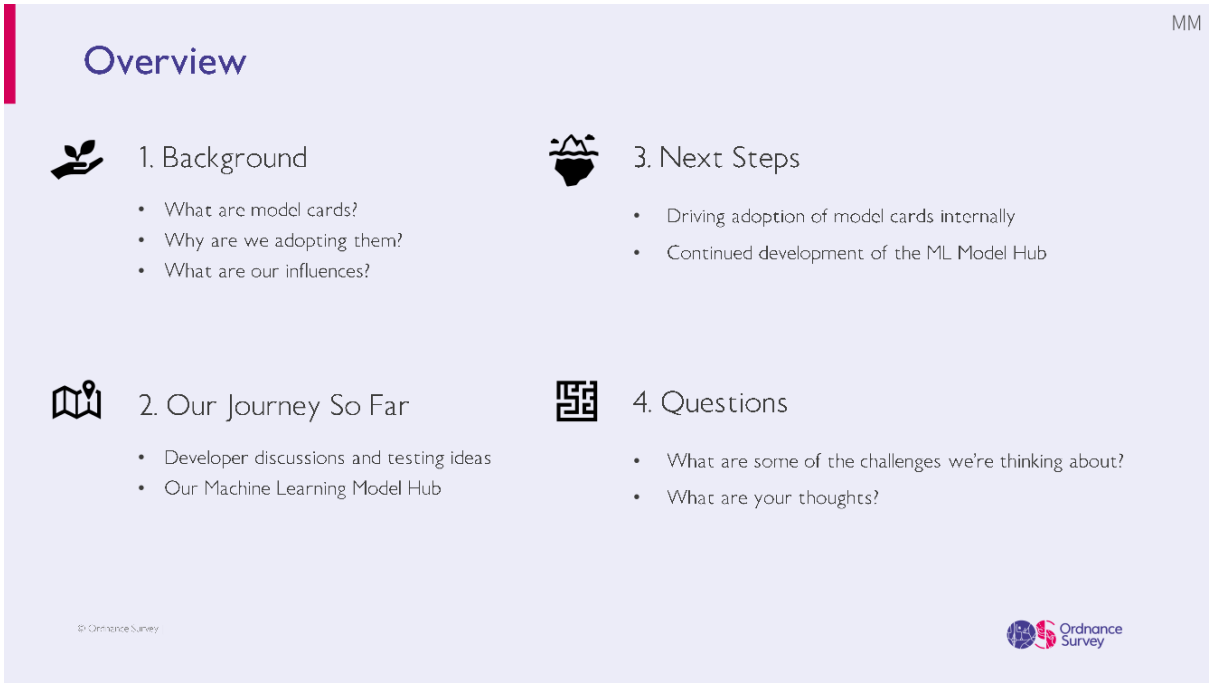

MM

Background

- OS uses machine learning alongside geospatial data for a variety of applications and customer problems – from deriving basemaps to enhancing our existing data using aerial and satellite imagery.
- We needed practical tools to encourage transparency and accountability in the creation of our machine learning models.
- That's where model cards come in – **they're a form of documentation which live alongside an AI model and tell you key information about it, such as what it was made for, how it was trained, how well it performs, how it should be used by others, and its potential biases and limitations**...

JW

Model Cards for Model Reporting

Margaret Mitchell, Simone Wu, Andrew Zaldivar, Parker Barnes, Lucy Vasserman, Ben Hutchinson, Elena Spitzer, Inioluwa Deborah Raji, Timnit Gebru
{mmitchellai,simonewu,andrewzaldivar,parkerbarnes,lucyvasserman,benhutch,espitzer,tgebruj@google.com
deborah.raji@mail.utoronto.ca

https://doi.org/10.1145/3287560.3287596

ARTICLE OPEN ACCESS

On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?

Authors: Emily M. Bender, Timnit Gebru, Angelina McMillan-Major, Shmargaret Shmitchell Authors Info & Claims

FAccT '21: Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency • March 2021 • Pages 610–623 • https://doi.org/10.1145/3442188.3445922

Ordnance Survey

MM

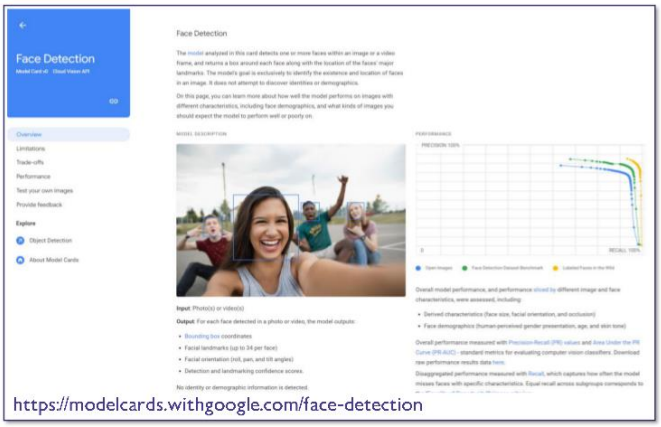

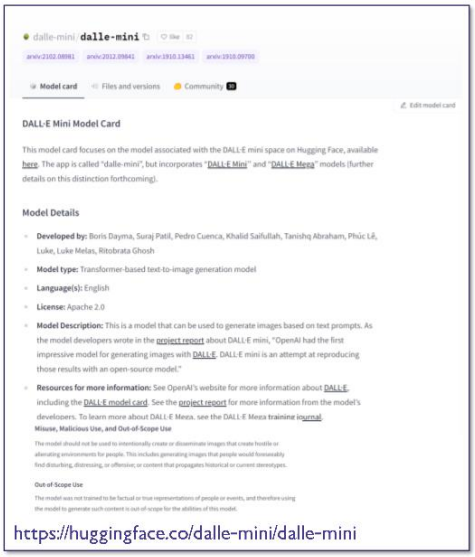

Ordnance Survey
SEE > BETTER PLACE

JW

Our Model Card Iterations

- Keeping a static model card in a model's code repository
- Using a template to automate a model card for each experiment
- Adding KwH / Co2eq calculations for sets of experiments
- Moving to a dedicated git repository per model card
- Publishing them to the web with a static site generator

Our Model Card Intentions

- Insight into what other teams are working on
- For developers, team leads and non-technical stakeholders.
- Facilitate transparency and reduce duplication of effort
- Drive data documentation for long-term maintenance
- Drive introspection about geodemographic bias
- In the future, link to models that have augmented OS products

Ordnance Survey

ML Model Hub Frontend

MM

https://huggingface.co/models

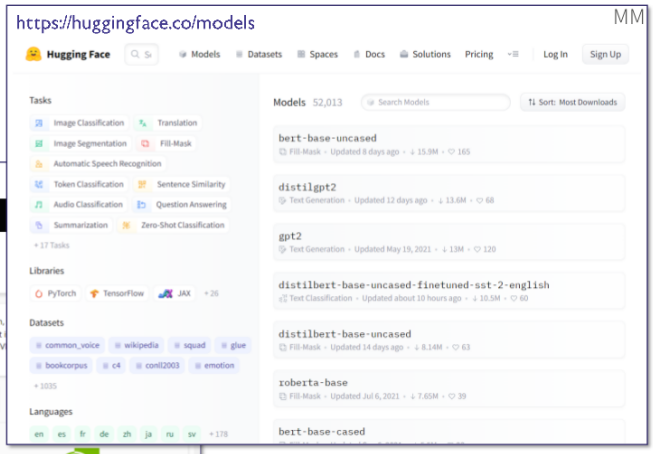

https://catalog.ngc.nvidia.com/models

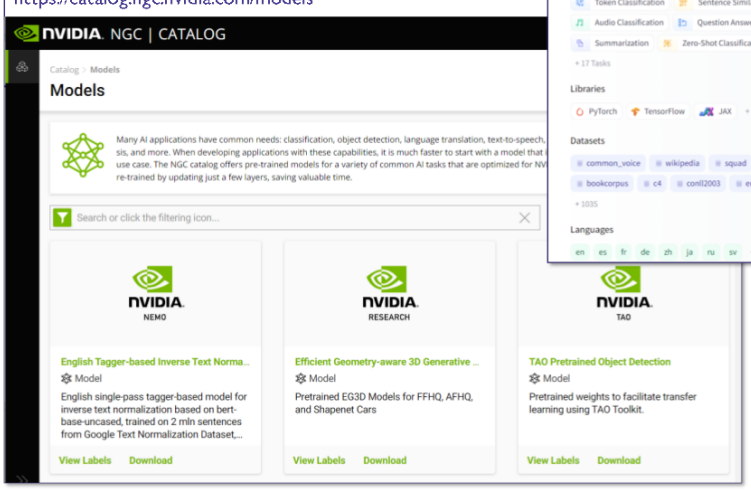

https://github.com/ivylee/model-cards-and-datasheets

Ordnance Survey

MM

DEMO: ML MODEL HUB FRONTEND

Ordnance Survey

JW

Questions

- Keeping them up to date
- Including them in the whole lifecycle
- Applying the mandate in a bottom-up context
- Treatment of production and experimental models
- How can we use it to drive transparency? City of Helsinki AI Register –

Ordnance Survey

JW

THANK YOU

We have been:

Mel Marochov Mel.Marochov@os.uk - OS Rapid Prototyping Team

Jo Walsh Jo.Walsh@os.uk - OS VisionAI Team

Ordnance Survey

D. Slides presented by Nigel Edmead & Vivek Sakhrani – Hyperlocal Mapping for Positive Change

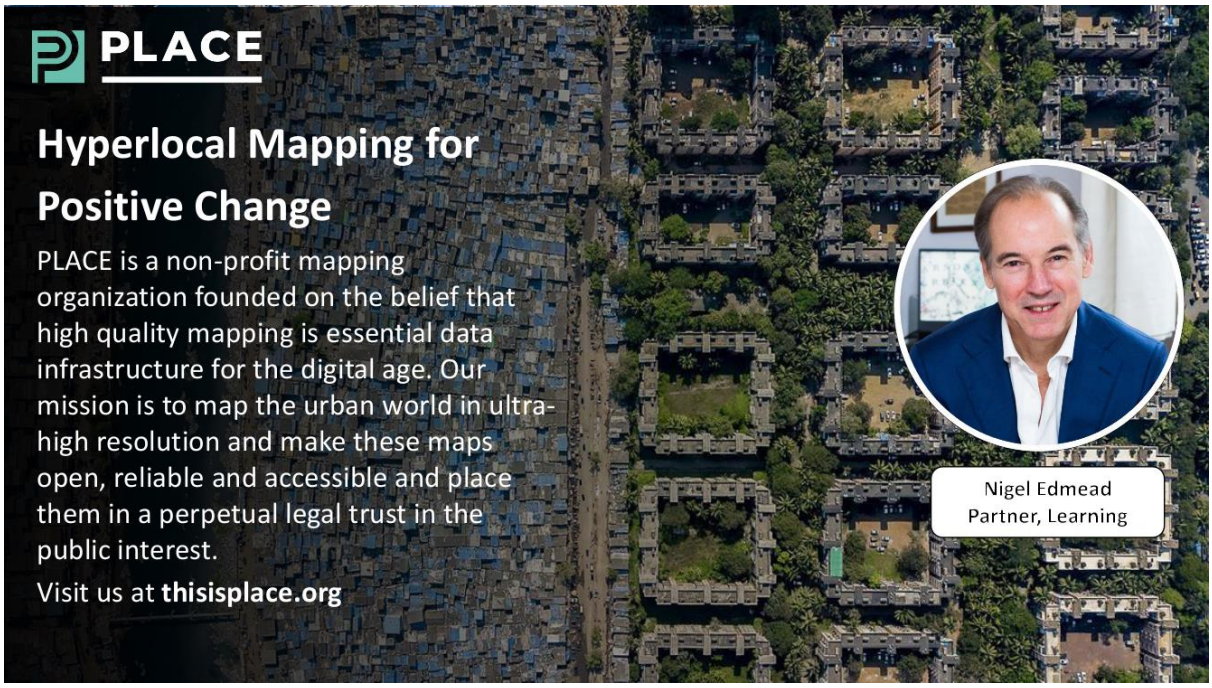

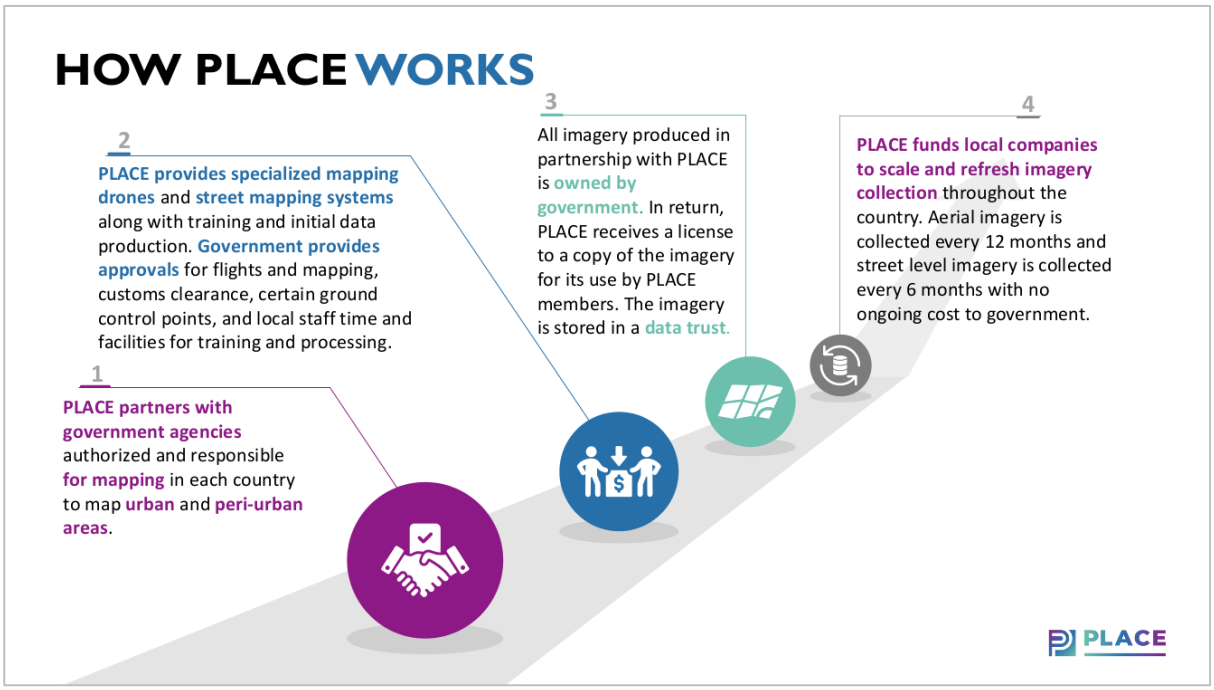

PLACE IMAGERY

PLACE only collects imagery so that members (like AtlasAI) can create maps and other value-added products and services.

- Geo-referenced aerial imagery
- 5cm GSD resolution
- RGB color

Delivered as an image collection. Images are geotagged by means of carrier phase differential GNSS and exposure event marking yielding fixed and float solutions.

- Geo-referenced terrestrial imagery
- High resolution
- RGB color

Delivered as an anonymized (face and number plates) image collection. Images are geotagged by means on board GPS.

PLACE

PLACE TRUST

PLACE stewards the imagery in a permanent legal data trust based in the United Kingdom

- A data trust provides independent, fiduciary stewardship of data.
- All imagery produced in partnership with PLACE belongs to the government of each country.
- PLACE receives from each government an irrevocable, perpetual, royalty free license to a copy of all imagery and its use by PLACE members through the PLACE Trust.
- The Trust will issue licenses for use of this imagery by our members.

PLACE

MEMBERSHIP

Membership is required to access and use PLACE imagery

- PLACE Trust is open to all organizations that agree to membership terms and conditions of use and behavior
- Member agree to ethical use of data principles outlined in the Locus Charter and PLACE provides recourse for violations
- Recognizes legitimate national security and community at risk concerns
- Commercial license allows for ownership of derived work. Non-commercial Creative Commons type license is used to publish in public domain
- Membership fees are based on type and size of organization and used to sustain ongoing operations and data collection
- Members have a voice in how we are governed and ensure we stay true to our mission

PLACE

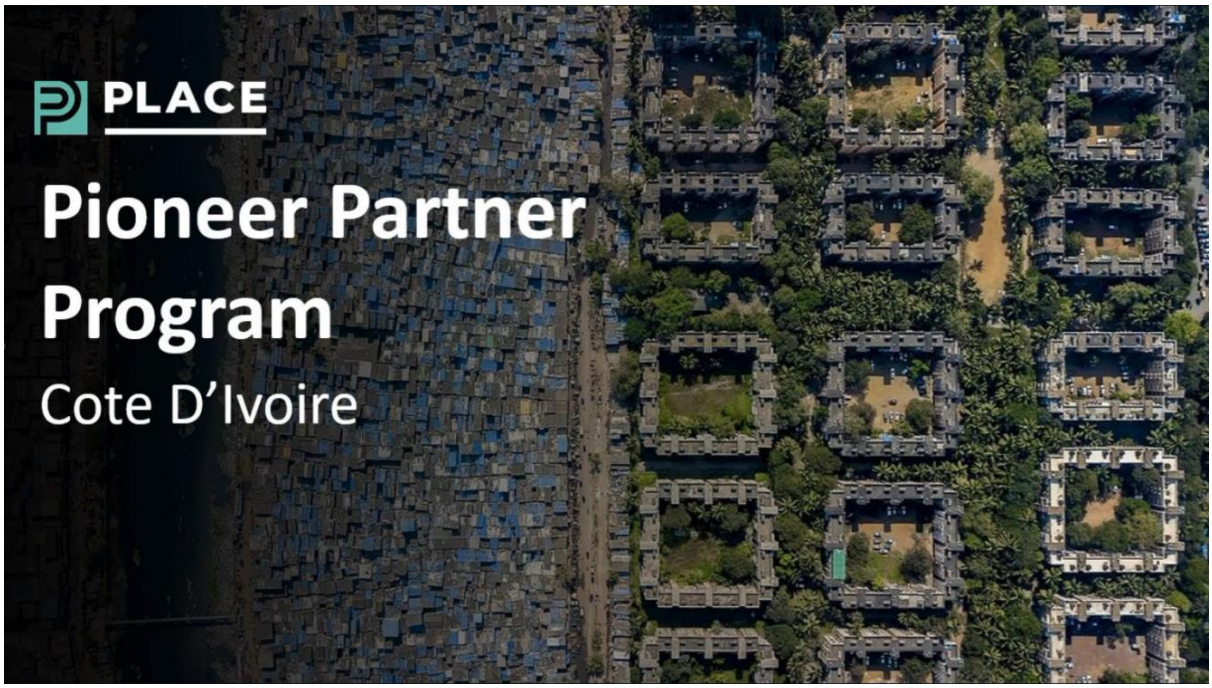

E. Slides presented by Ruth Bowyer – Ethics and public perceptions of georesearch: learnings from a population-based cohort

Ethics & public perceptions of georesearch: learnings from a population-based cohort

Dr Ruth C E Bowyer
King's College London & The Alan Turing Institute

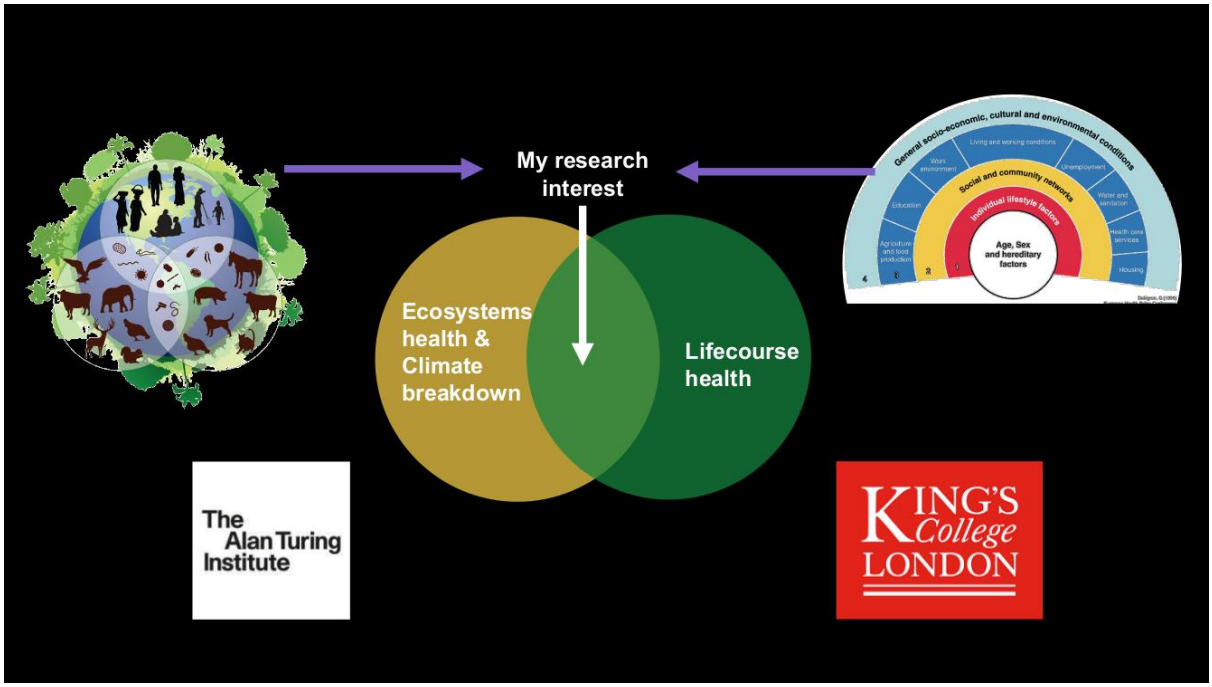

We should involve people in research that affects them

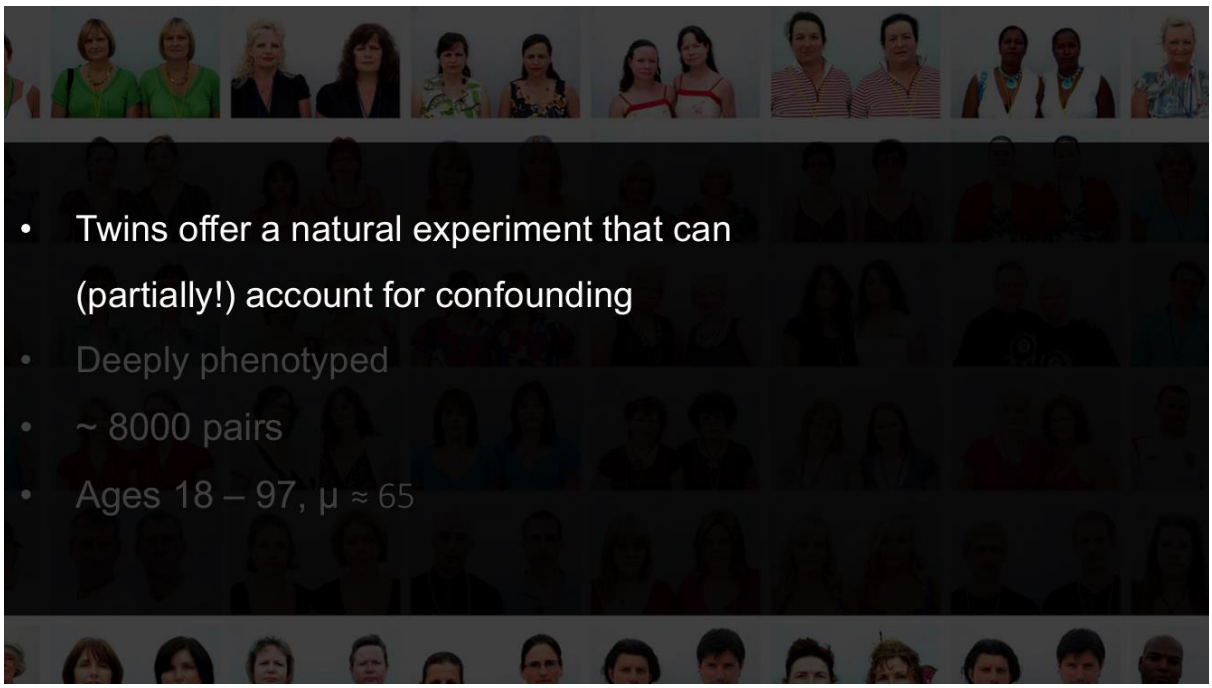

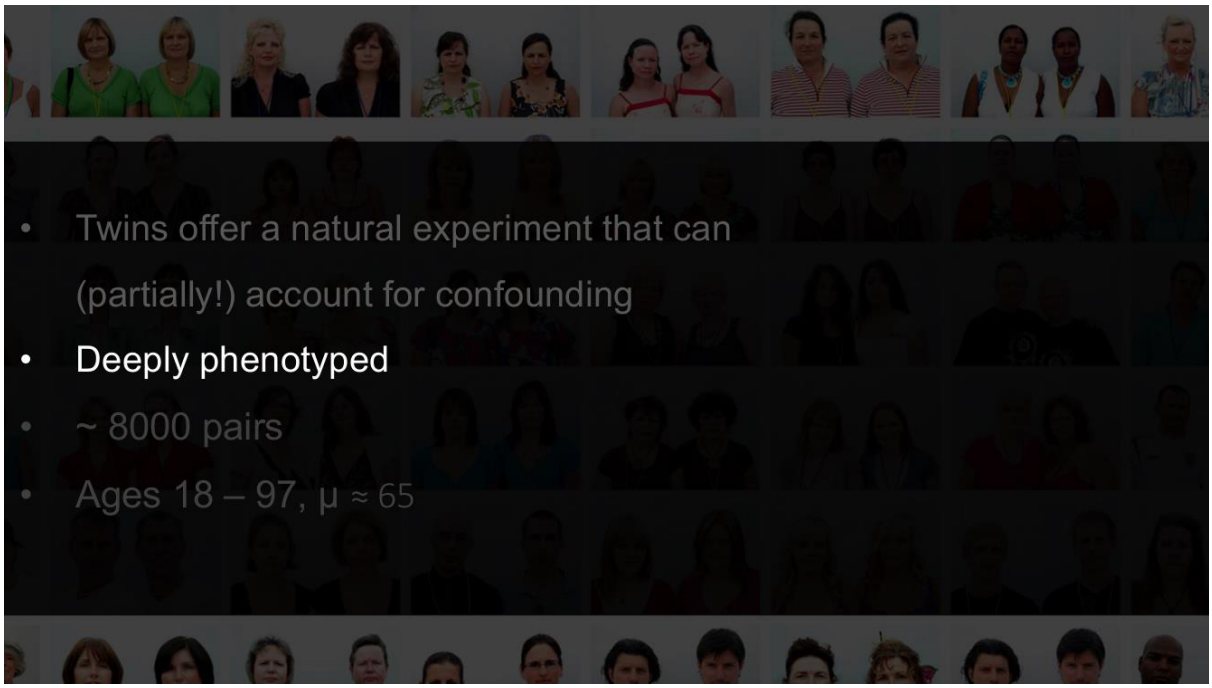

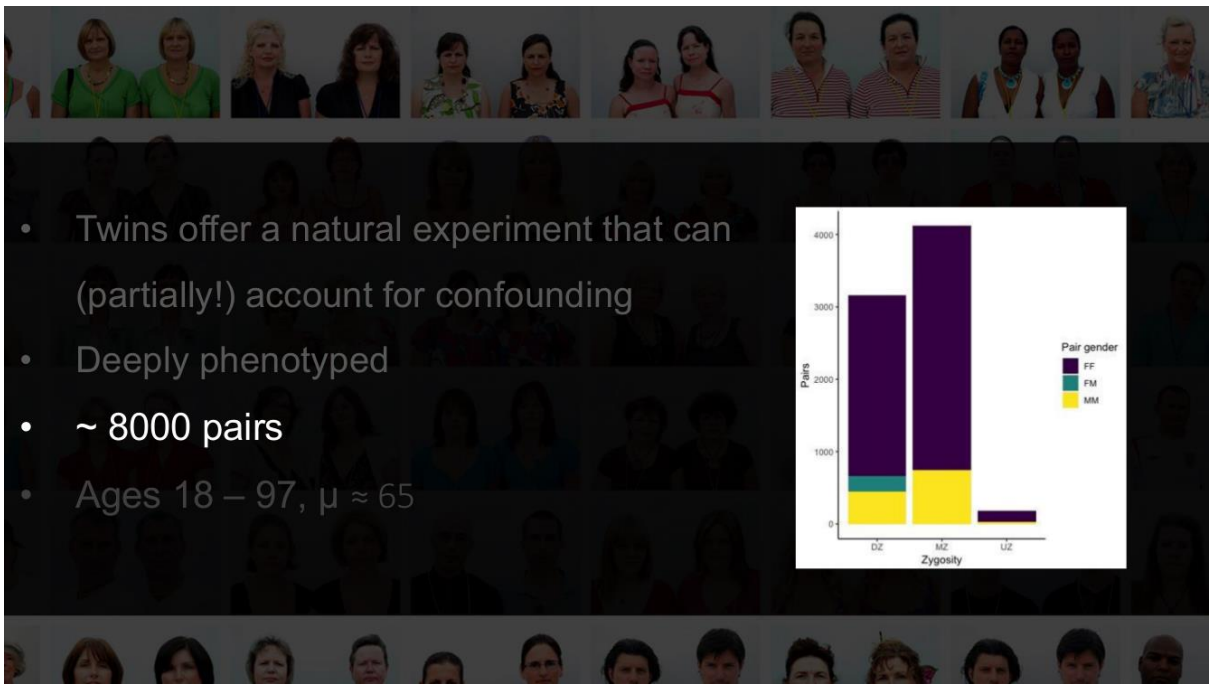

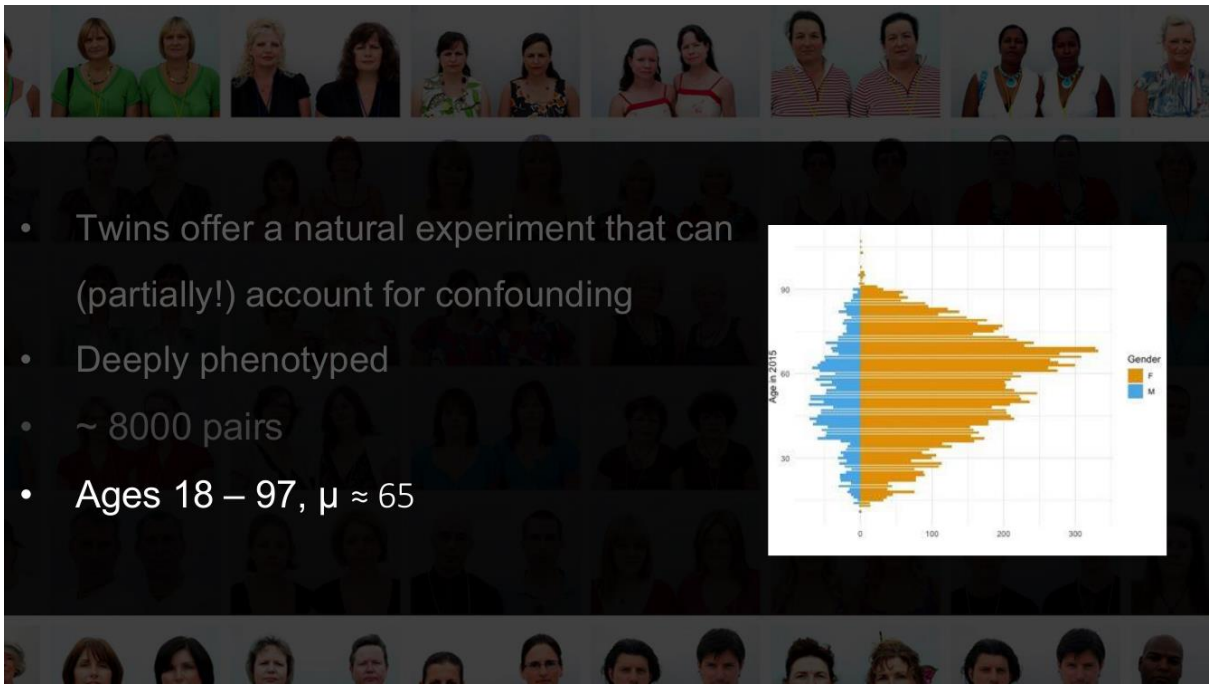

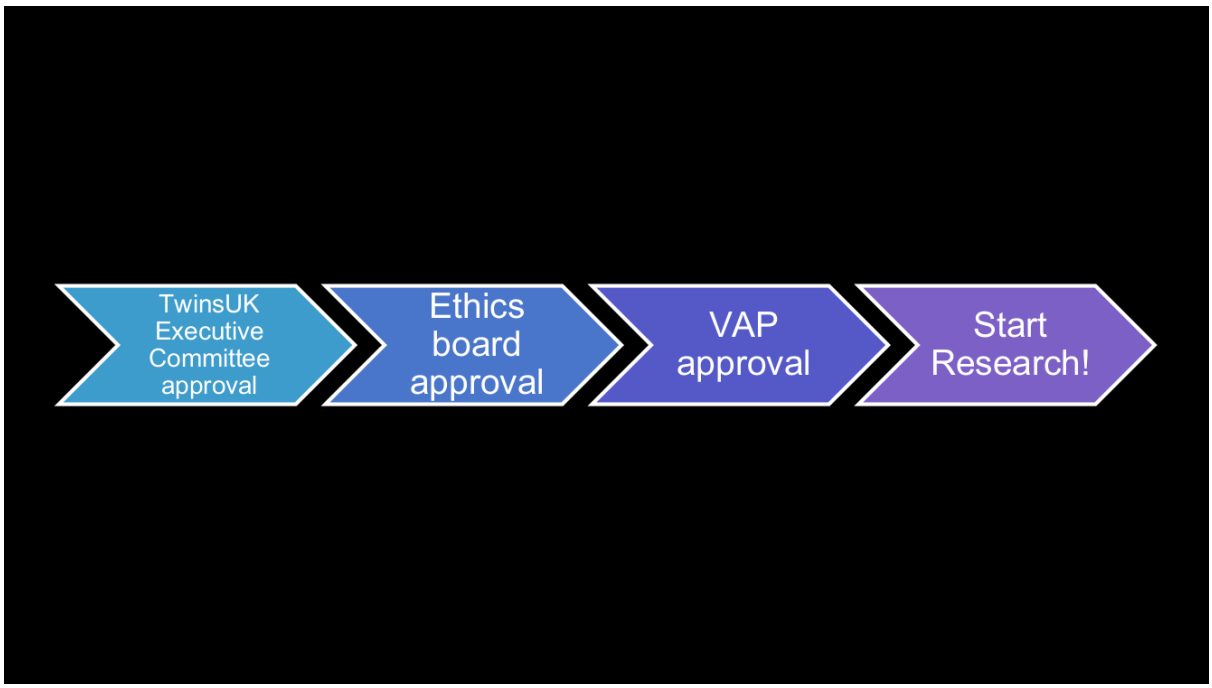

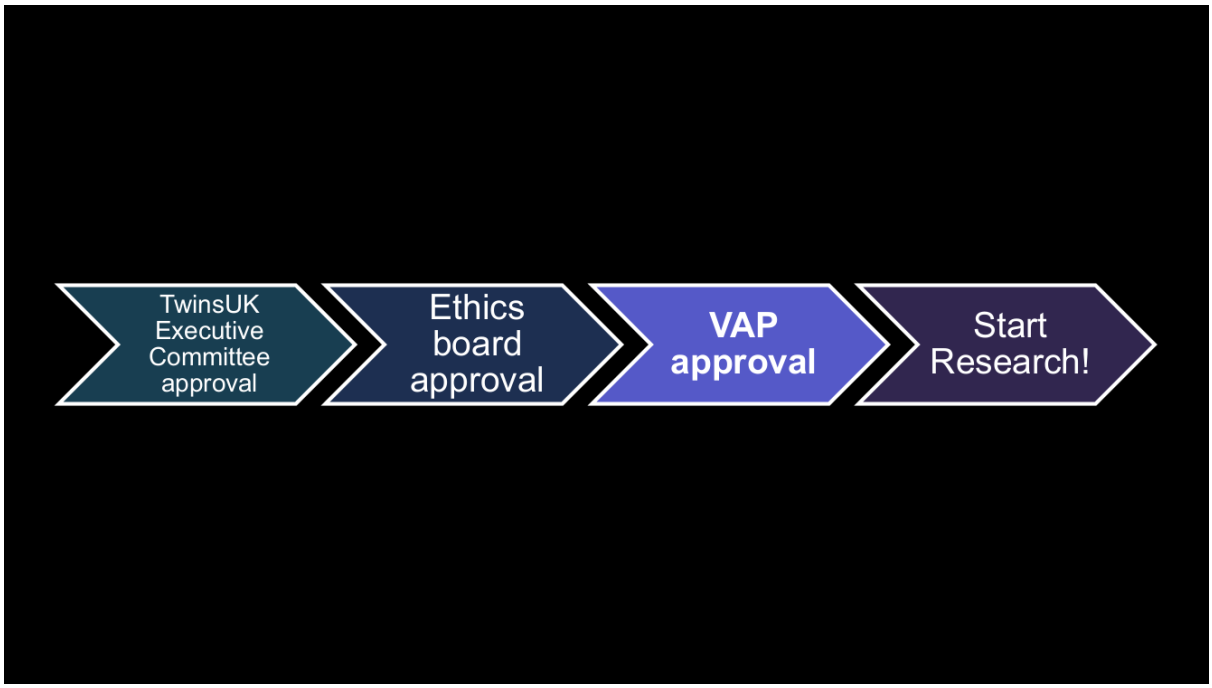

Case study: Mapping the envirome in Twins

Understand the contribution of environmental exposures to age-related decline and resilience

How do you feel about wearing a GPS tracker?

- 'Google already does this'
- Preferred that their data be used for health research
- The ability to turn off trackers was appreciated
- Participants wanted measures to be further anonymized when sharing with researchers external to TwinsUK

How do you feel about studies associating your local environment with potential negative health outcomes?

- 'Media organisations already do this with best/worst lists'
- Participants agreed they would be very interested to find out more about how local conditions were impacting their health
- Potential for empowerment in engaging with local councils

Involving participants in the research design

- Challenged assumptions
- Improved the research design
- Empowered participants

We should involve people in research that affects them

Thank you!

ruth.c.bowyer@kcl.ac.uk
rbowyer@turing.ac.uk

F. Slides presented by Jeni Tennison – Public participation

Three layers of participation

For organisations that steward data

- **policy** – setting the bounds of activity
- **operational** – making detailed decisions
- **accountability** – detecting when things go wrong

CONNECTED BY DATA

4

Geospatial Commission, **Public dialogue on location data ethics** (November 2021). https://www.gov.uk/government/publications/public-dialogue-on-location-data-ethics

CONNECTED BY DATA 5

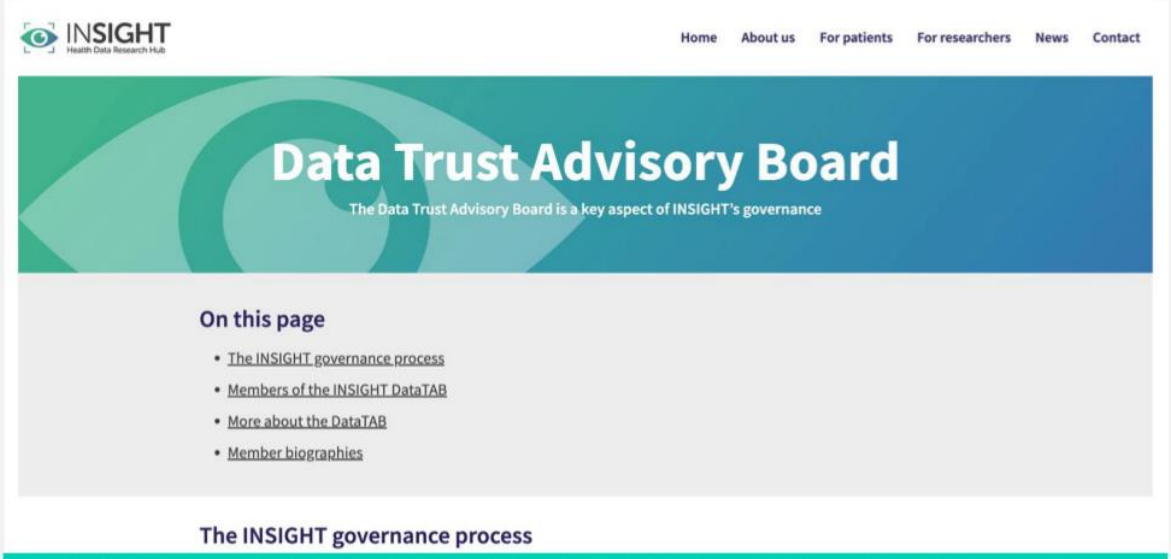

INSIGHT, **Data Trust Advisory Board** https://www.insight.hdrhub.org/data-trust-advisory-board

CONNECTED BY DATA 6

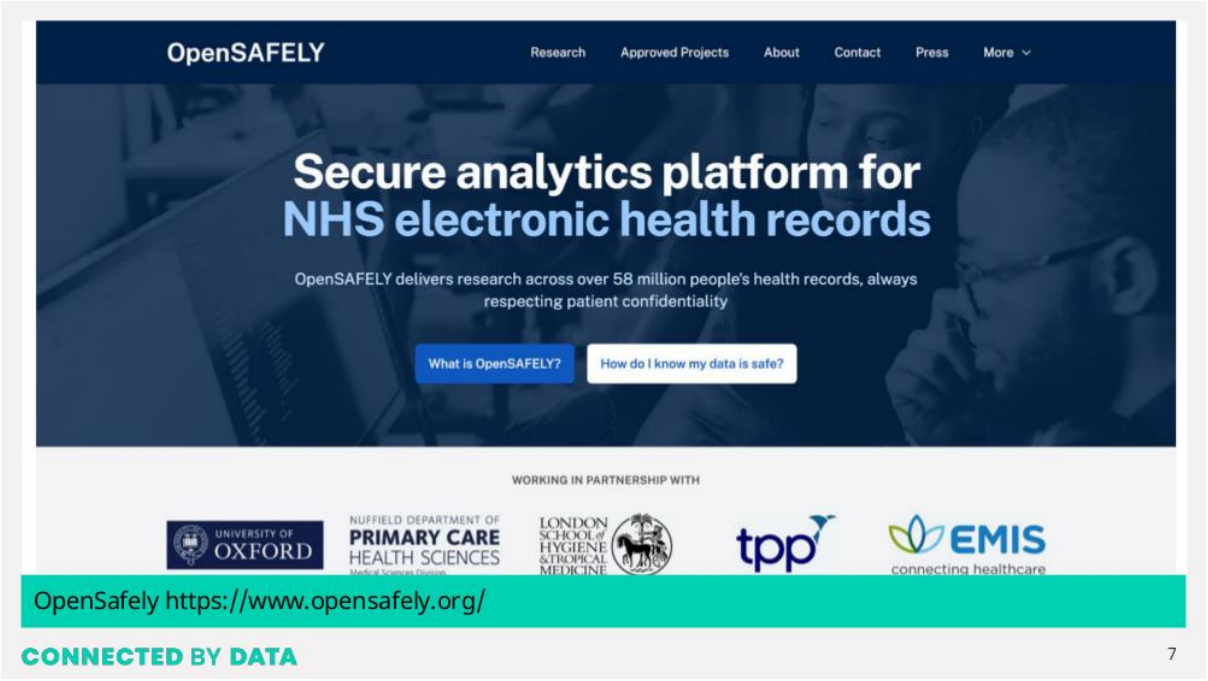

G. Slido surveys

The three surveys during the workshop helped us understand the motivations and perspectives of the participants. Around 30 people engaged with Slido (including the Q&A) and around 20 took part in the surveys (dropping off towards the end of the day).

Morning survey: Getting to know each other

What sector best defines your current role? Votes: 18

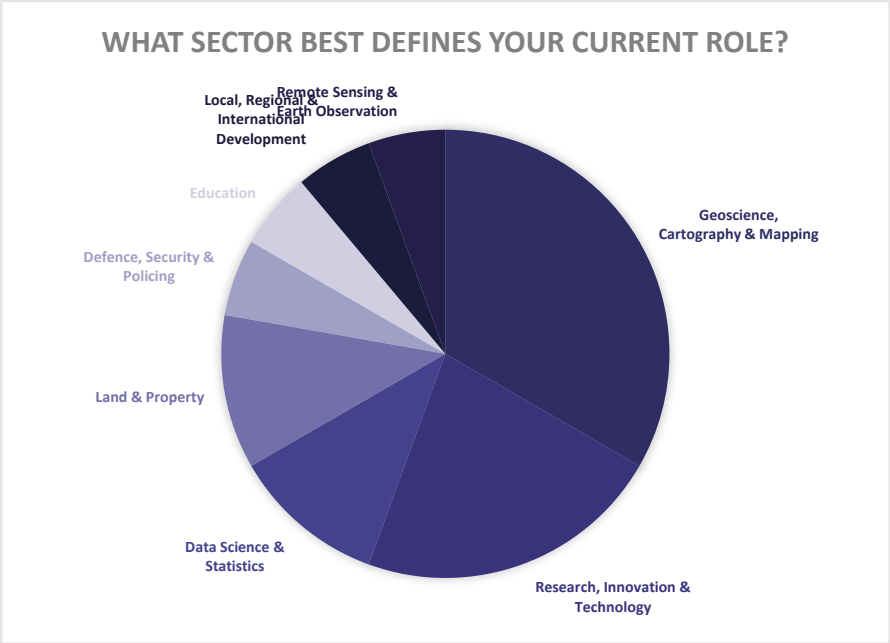

How well-informed do you feel about AI Ethics? Votes: 18

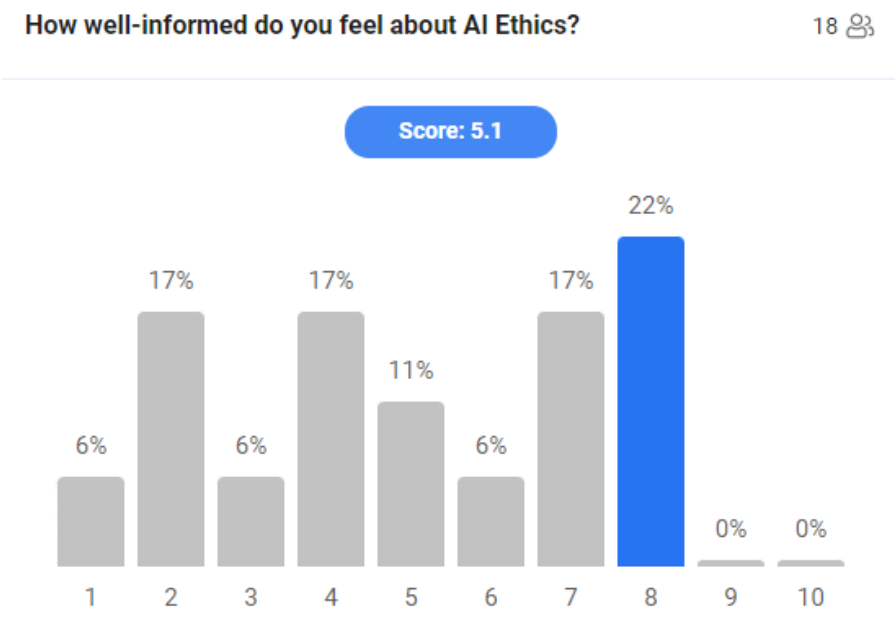

What motivates you to join this workshop? Votes: 17

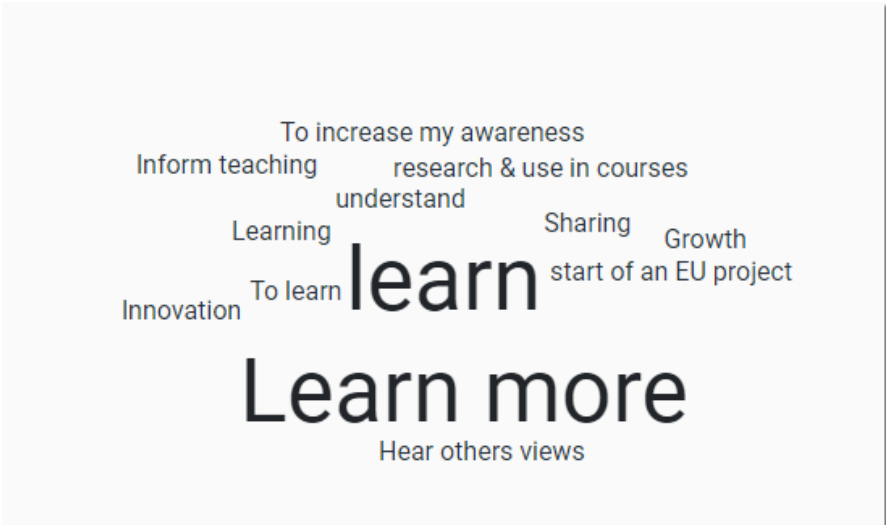

Lunchtime survey: Where do we find ourselves?

My organisation has explicitly subscribed to a set of ethical principles Votes: 16

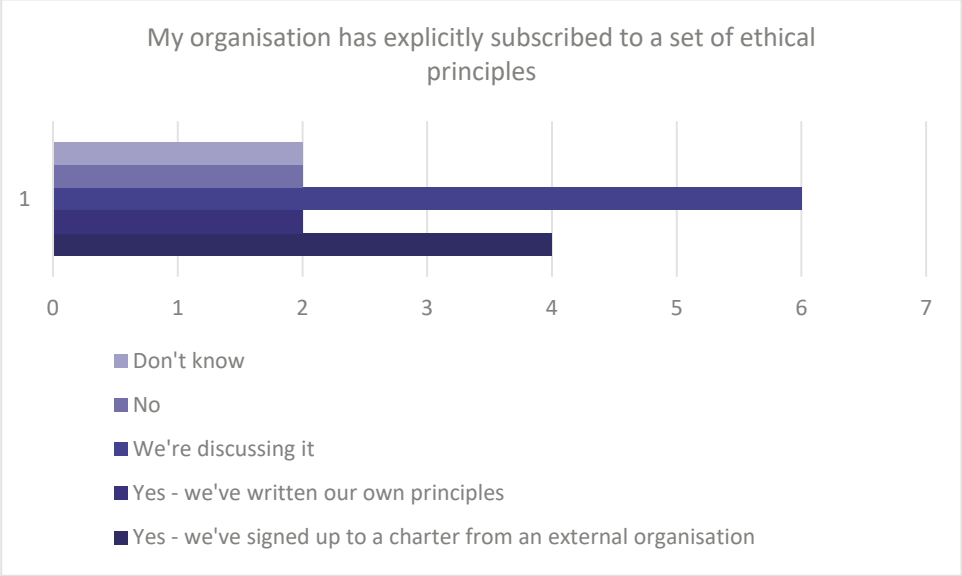

What ethical practice, if any, have you implemented within your work? Votes: 6

- Data Ethics Code of Conduct
- Supplier required us to report on sustainability, which forced us to think about this
- Inclusive Data Charter
- Ethics Review Committees
- Ethics assessment and ethics training when proposing research
- Requirement by some funders for an "Ethical Board"

Final survey: What will we take forward?

How well-informed do you feel about AI Ethics? Votes: 9

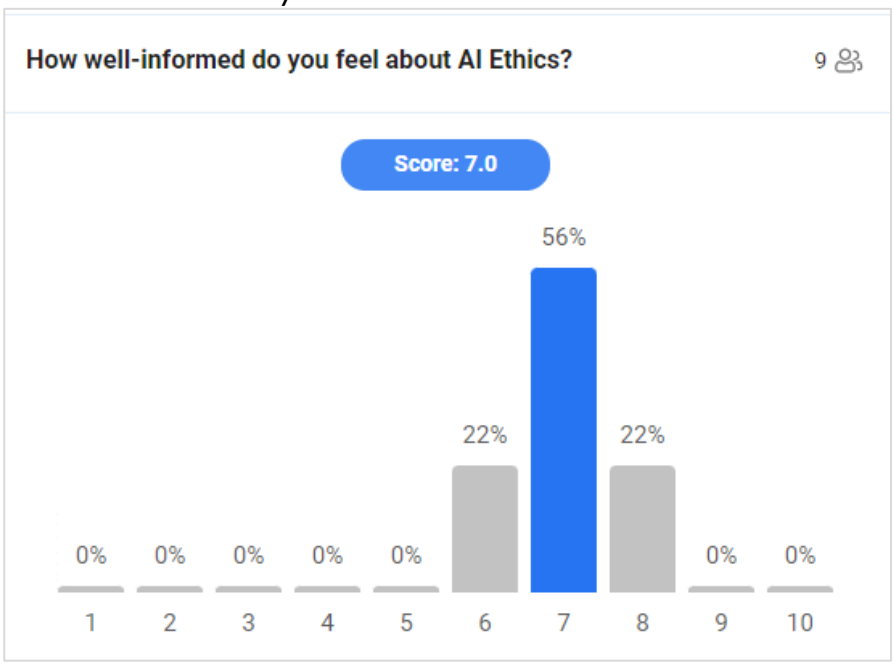

As was pointed out during the workshop the apparent improvement in how informed participants feel about AI Ethics is likely to be due to self-selection as the more-informed have stayed on to this stage.

What ethical practice would you like to have initiated or even implemented in your organisation in the next year? Votes: 6

- Consider the ordering and appropriateness of the Locus Charter principles at the start of each project
- Implementation of the Locus Charter in my organisation. To have a development programme based on supporting countries in implementing ethical data.
- Review of RGS and OGC materials Thinking about ethics in my work Thinking about how I include ethics in my training course
- All of them!
- Model cards
- Local involvement and Citizens assemblies

Tell us what is at the forefront of your mind at the end of this workshop Votes: 7

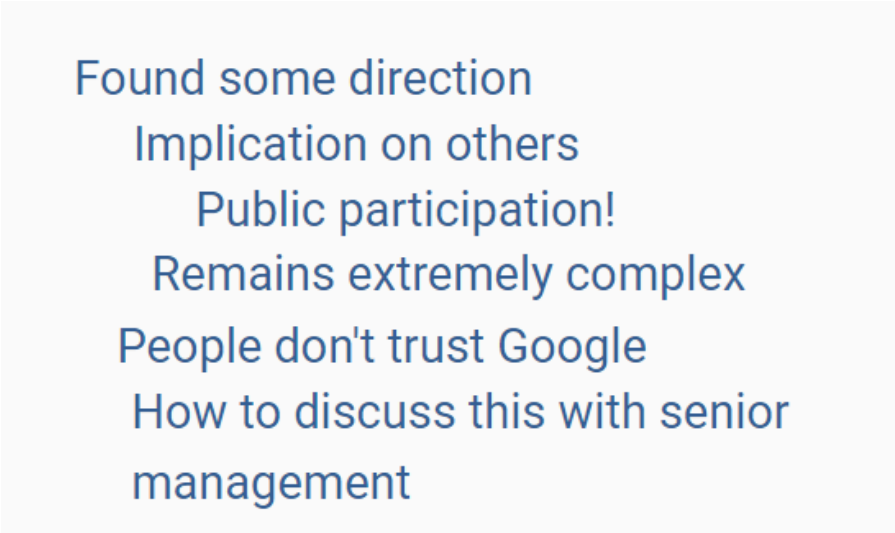

These results provide an essence of who attended the workshop and what was at the forefront of their thinking during the workshop.

H. Practices that identify and mitigate harms outputs

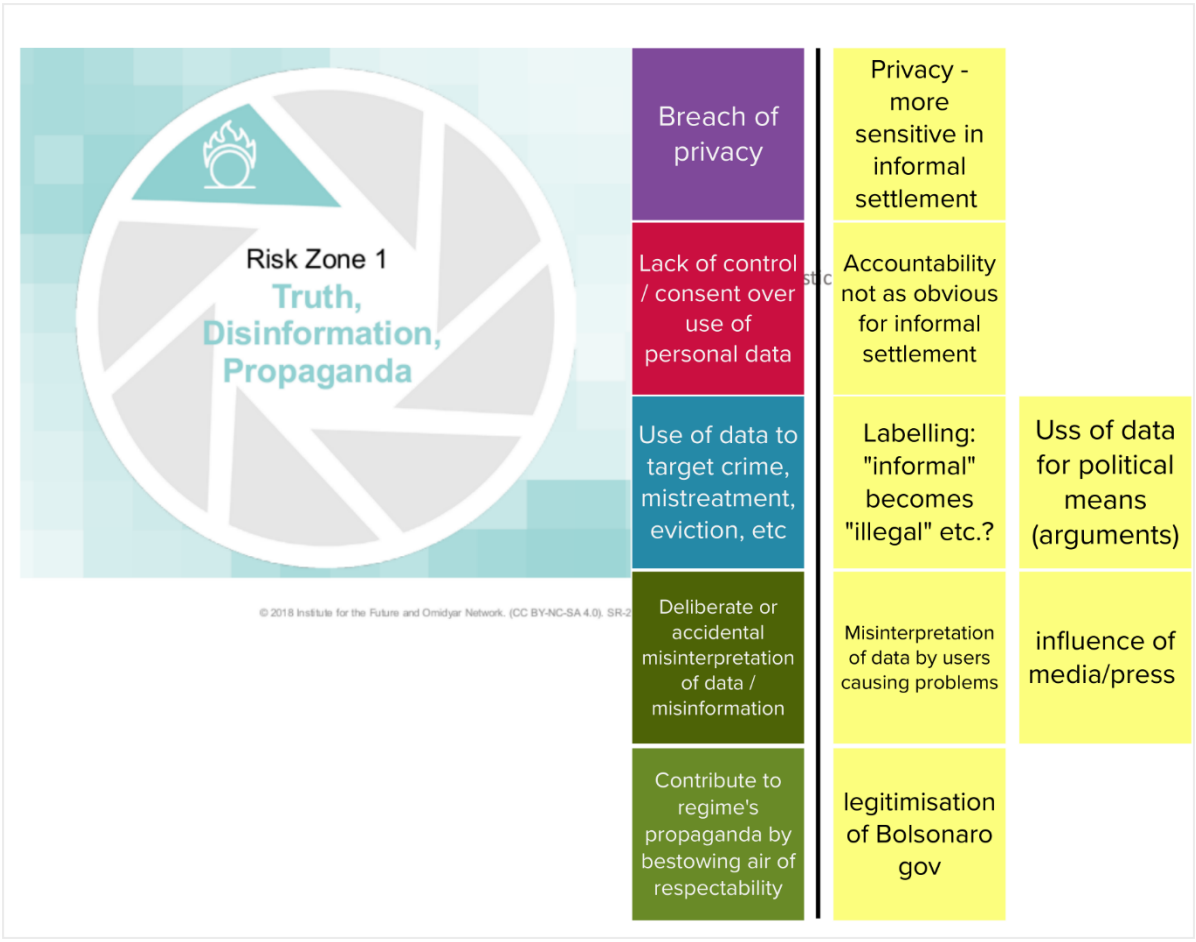

Figure H1: Harms that participants identified (yellow) and their higher-level category that related to Ethical OS Risk Zone 1 Truth, Disinformation, Propaganda

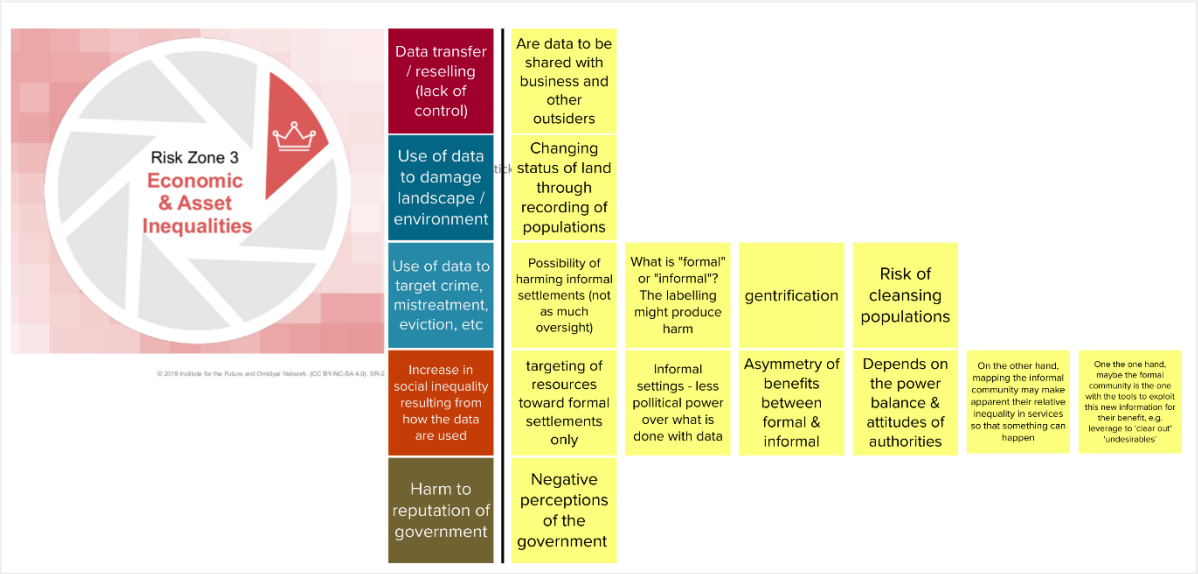

Figure H2: Harms that participants identified (yellow) and their higher-level category that related to Ethical OS Risk Zone 3 Economic & Asset Inequalities

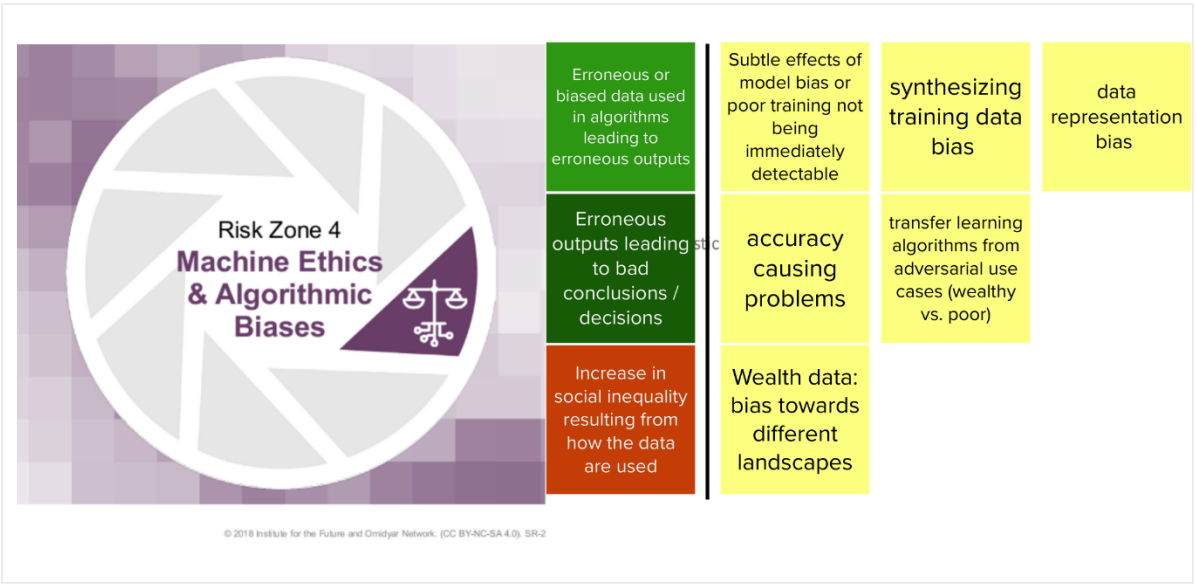

Figure H3: Harms that participants identified (yellow) and their higher-level category that related to Ethical OS Risk Zone 4 Machine Ethics & Algorithmic Biases

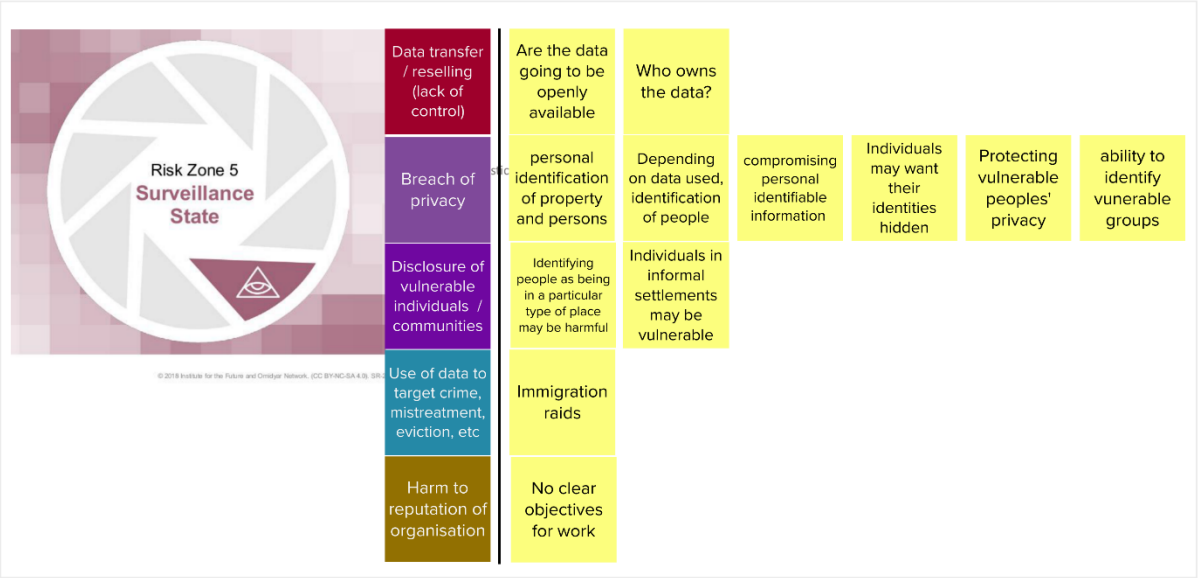

Figure H4: Harms that participants identified (yellow) and their higher-level category that related to Ethical OS Risk Zone 5 Surveillance State

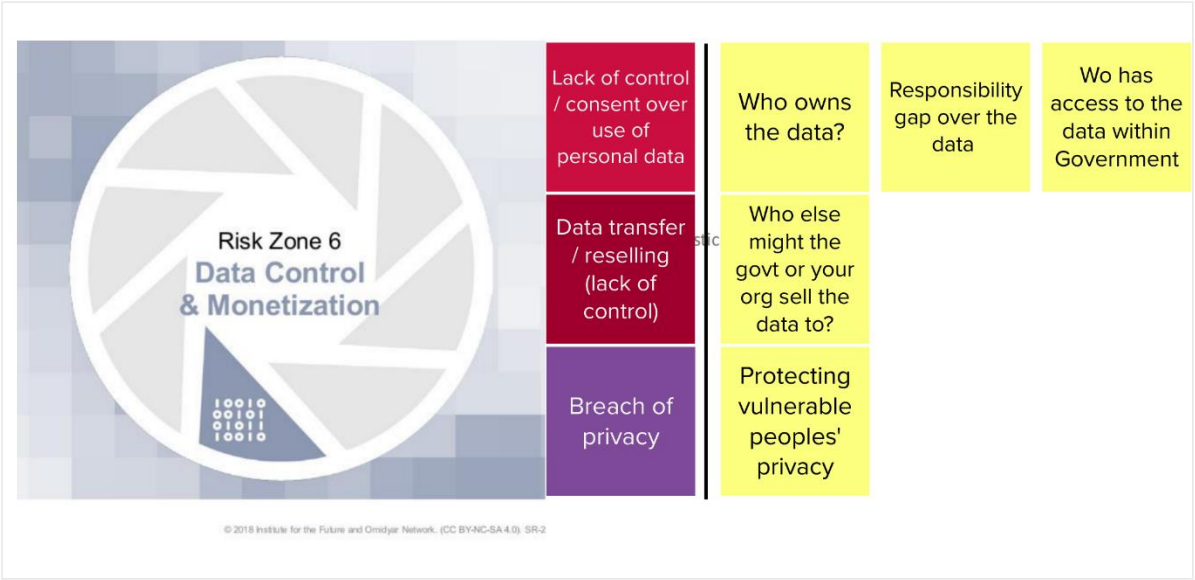

Figure H5: Harms that participants identified (yellow) and their higher-level category that related to Ethical OS Risk Zone 6 Data Control & Monetization

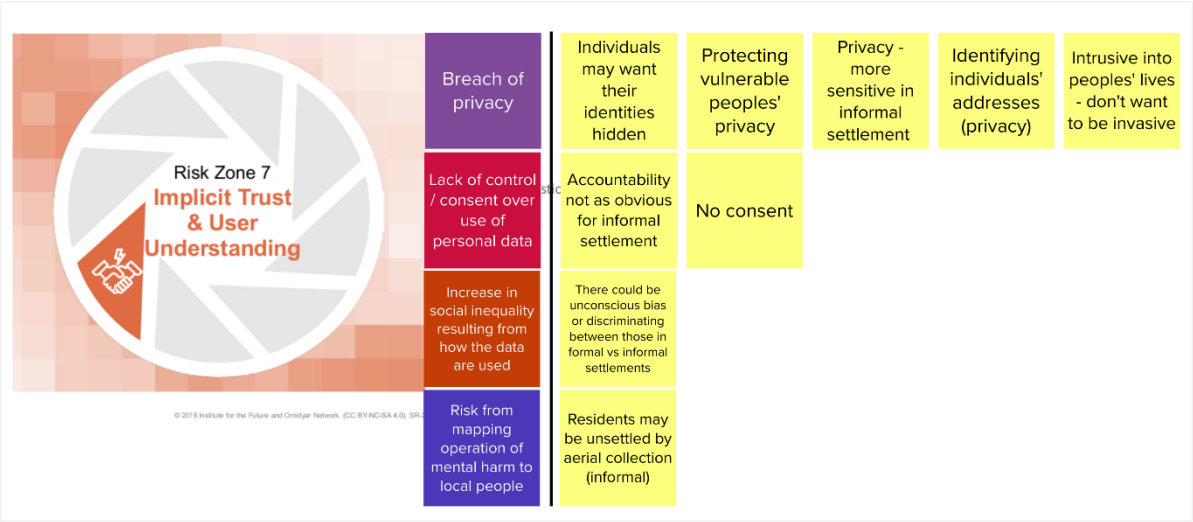

Figure H6: Harms that participants identified (yellow) and their higher-level category that related to Ethical OS Risk Zone 7 Implicit Trust & User Understanding

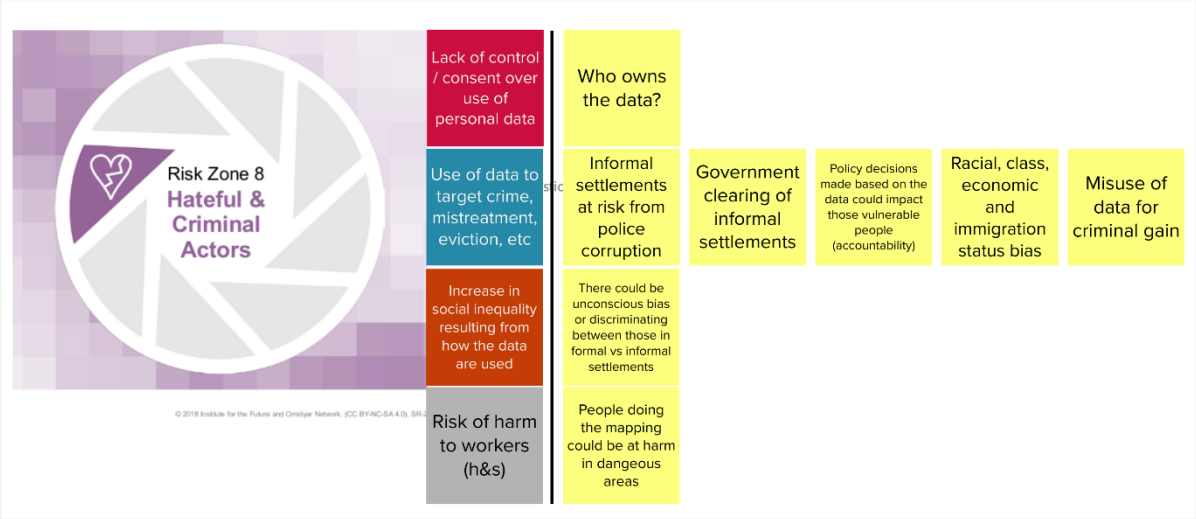

Figure H7: Harms that participants identified (yellow) and their higher-level category that related to Ethical OS Risk Zone 8 Hateful & Criminal Actors. Risk of harm to workers should be covered under Health & Safety assessments.

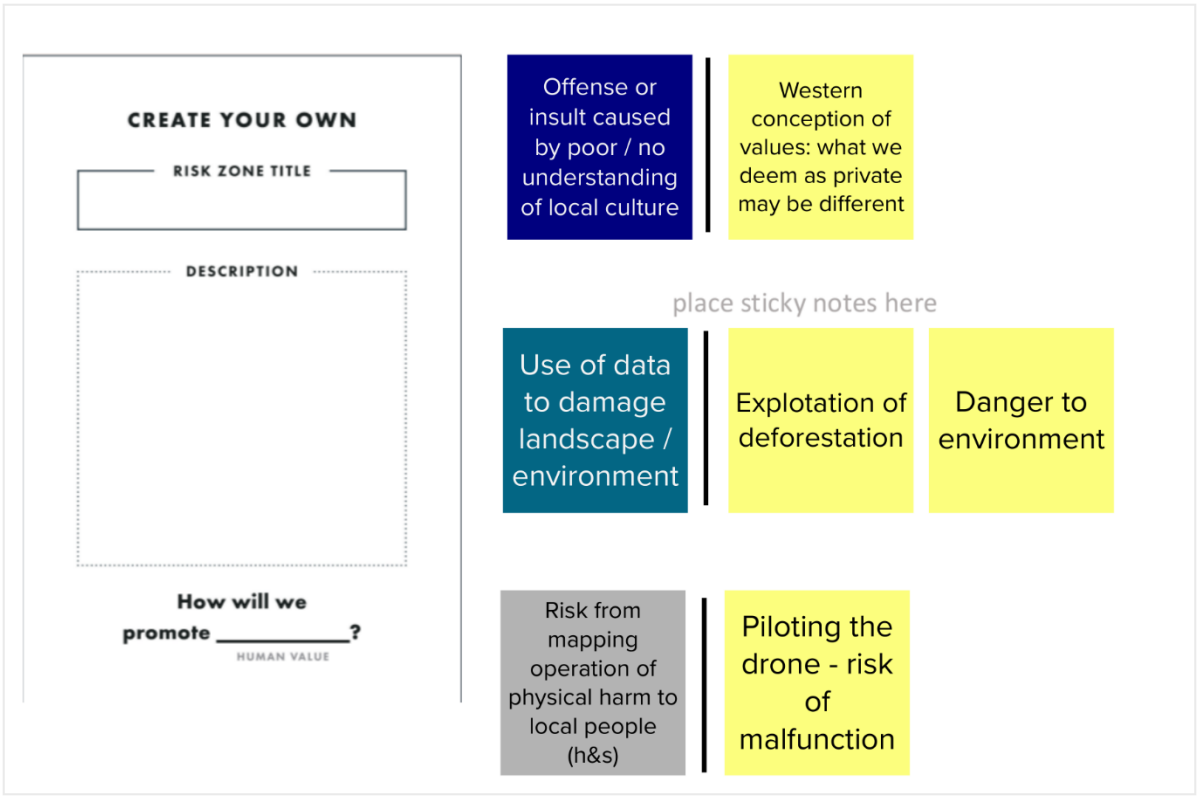

Figure H8: Harms that participants identified (yellow) and their higher-level category that fell outside of the Ethical OS Risk Zones. Risk from mapping operation of physical harm to local people should be covered under Health & Safety assessments.

Step 3: Identify practices that would ensure these harms were identified and / or mitigated - 30 minutes

For each of your identified harms, collectively decide what organisational practises could be put in place to identify, avoid or mitigate the harm.

If you need a little inspiration, there are resources in the GeoAI Ethics workshop resources board. You can visit that board by clicking here.

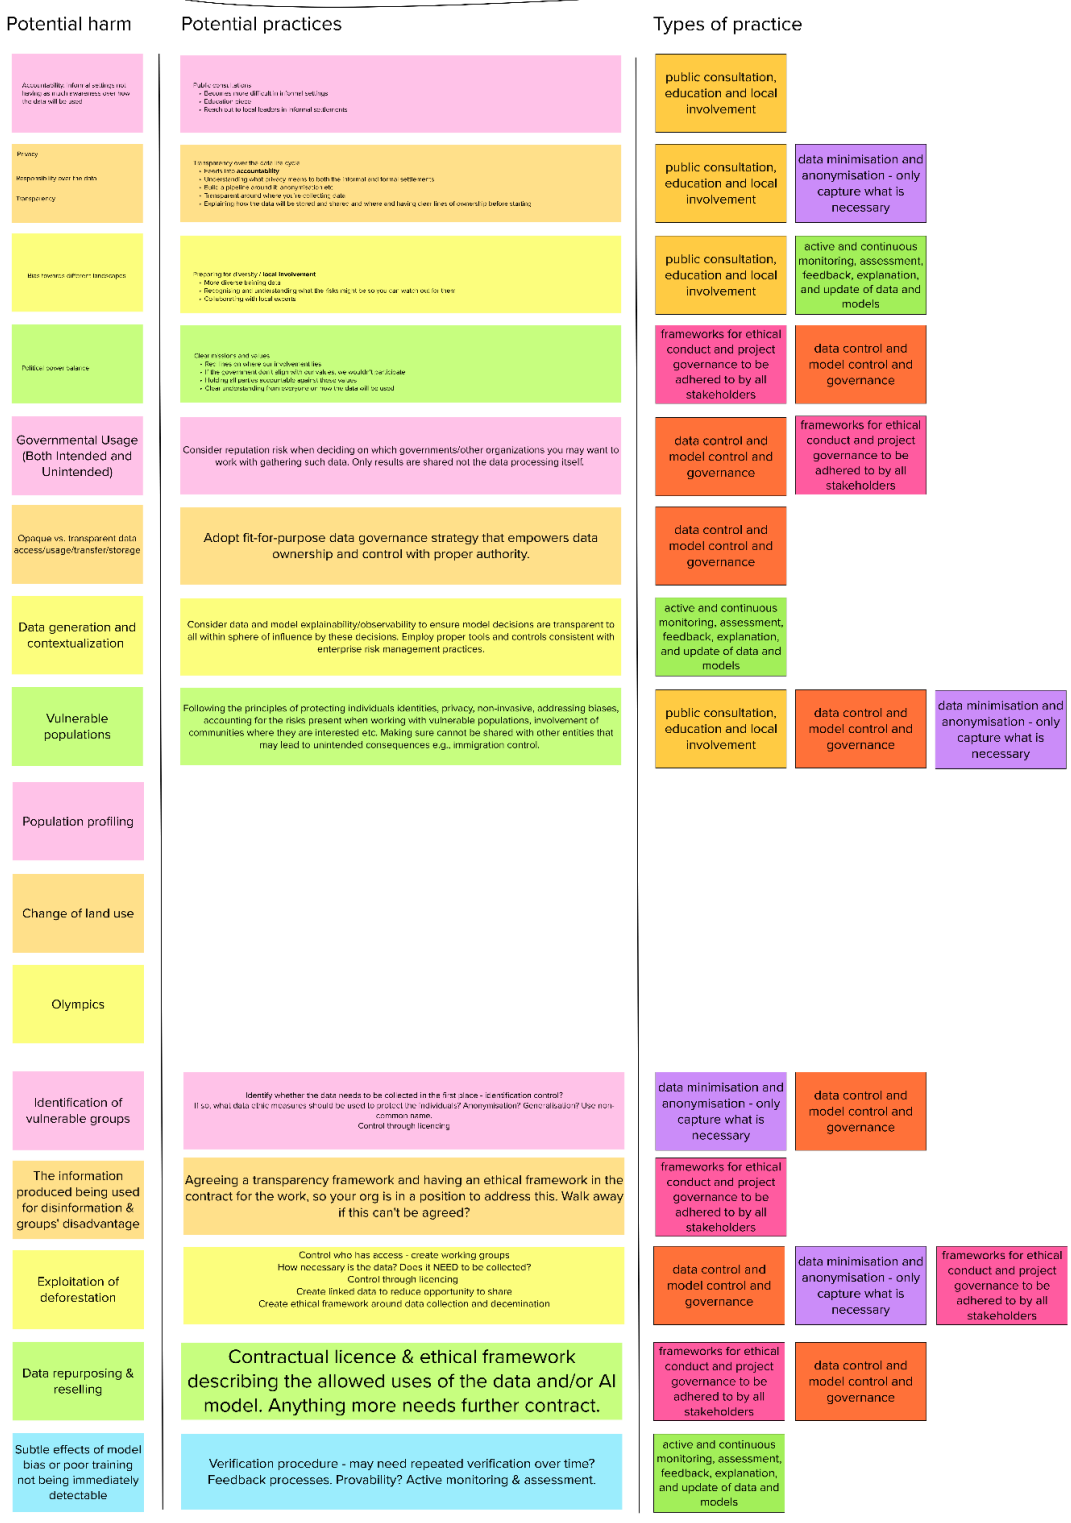

Figure H9: Consolidation of suggested practices from all the breakout groups and summary into 5 different types of practice