Chapter 2
Ethical and legal concerns on data science for large scale human mobility

Albert Ali Salah[1,2], Cansu Canca[3,4], Barış Erman[5]

2.1 Introduction

Big data based analysis of human mobility relies on various data sources and combines these to address a range of "wicked" social problems—that is, complex problems without a clear analytical solution and with many dimensions that need to be optimised simultaneously. These problems may not have clear cut causal relationships, but as in most complex systems, they are driven by feedback loops, circular causality, and parameters that involve many stakeholders. A purely data-driven approach could be dangerous, simply because it neglects to address this complexity properly and thereby generates various side issues while only solving the problem in focus.

This chapter introduces the ethical and legal aspects of data science, with the aim of providing the reader with a basic understanding of the ethical risks, legal issues, and tools to mitigate them. These risks may ensue from the use of a particular approach or methodology, or they may have deep-seated causes, such as problems in key definitions, issues related to specific data sources, and issues in policy recommendations.

Section 2.2 introduces applied ethics and ethics principles in the context of data science. Section 2.3 details some common ethical issues in data science and technologies that build on data science, describing how these issues may manifest in the migration and mobility domains. Section 2.4 provides a list of ethics tools for projects and initiatives, while Section 2.5 briefly addresses legal concepts and concerns, which are more closely related to policy makers and end users but still relevant to researchers. Further reading is provided at the end of the chapter.

[1]Utrecht University, the Netherlands; [2]Boğaziçi University, Turkey; [3] Northeastern University, USA;[4]AI Ethics Lab, USA; [5]Yeditepe University, Turkey.

2.2 Applying ethics and ethics principles

2.2.1 What is applied ethics?

When we discuss the ethics of a particular technology, our goal is to determine how we can ethically develop and use that technology in various domains. This requires an understanding of the ethical risks already entailed in the proposed technology and existent in the domain within which this technology is meant to be utilised. Through an ethics analysis that explicates how these risks may manifest as the technology enters into use in a given the domain, we can determine the best course of action to mitigate ethical risks. In this chapter, we are focusing on the ethical questions and concerns related to data science technologies as they apply to the domain of migration and mobility.

Ethics, more specifically *applied ethics*, aims to guide actions. Applied ethics is a normative subdiscipline of philosophy, where the core question is "what is the *right* thing to do in a given situation". Here, "right" is understood in terms of *good* and *just*. This might be, for example, about a developer choosing which dataset or algorithm to use in order to minimise harm or a policymaker deciding which safeguards to set up in order to reduce algorithmic discrimination. In order to guide such decisions, an ethics analysis would employ theories from moral and political philosophy, engage with morally relevant concepts, and use analogies and thought experiments to test the argument. Since the questions at hand are real-world questions, such an analysis would also require information provided by other disciplines. For example, if the question is whether it is ethical to mandate a stricter border control during a pandemic, we need information regarding the transmission mode and rate of the virus, the mortality and morbidity rate of the disease, the expected social and economic impact of the policy on various groups, the feasibility and cost of this policy, and details of the other available options, just to name a few. Once this information is provided by health sciences and social sciences, we can analyse the ethical justifications and implications of this question in terms of harm and benefit, individual freedom and autonomy, and distribution of benefits and burdens within the society to determine whether the policy is ethically permissible, necessary, or ethically prohibited.

2.2.2 Ethics in data science

When it comes to data science for migration and mobility, ethical questions arise from the domain of migration itself, from the types of technologies that can be utilised or developed in relation to migration, and from the intersection of these two. These questions are present from the conception of a project to the presentation of the result, including ethically loaded decisions in formulating the problem and choosing the methods and tools for researching it.

Concerns in data ethics can be distinguished as those related to data, algorithms, and practices (Floridi and Taddeo, 2016). Ethics of data pertain to the collection, storage, and usage of large scale data, re-identification and privacy issues, consent of data owners, the biases inherent in the dataset itself, and risks and benefits arising from the analysis of the data. Ethics of algorithms focus on the complexities of algorithms, designers' and developers' responsibilities for ethical design, and auditing and transparency. Finally, ethics of practice relate to the practice of data science and the results of data analysis and actual deployment of algorithms, their effects on real-world decisions, focusing on power, authority, policy, and user rights. These areas are necessarily intertwined, and the boundaries are not clearly distinguishable.

For example, a big data project investigating migrants' access to schools and educational facilities should take into account potential issues in the formulation of the project and the use of different data sources. This may, for instance, include considerations like data coverage in rural areas to ensure there are no data gaps and the consent of data owners. Mamei et al. (2019) provide an example in this area that uses mobile call detail records (CDR) data to compute refugees' physical access to educational institutions. The algorithms used to estimate access should take specific biases into account (*e.g.* assumptions about modes of transport or considering gender-related issues in transportation). Once such an analysis is completed, the policy recommendations could take a broad range of possible consequences and complex factors into account, as the changing use of educational facilities may have different impacts on different populations. Clearly, considering all three levels of ethical concerns (i.e. data, algorithm, and practice) is necessary to fully assess and address the ethical risks and benefits. An example is the closing down of the temporary language centres in Turkey, in order to integrate Syrian children to Turkish primary schools, rather than teach them at the language centres. This was done to improve their social integration, and indeed served this aim. However, this policy had a strong negative effect on the older female refugee population, who were able to attend these centres but not able to attend the primary schools to learn the local language (Haznedar et al., 2018; Salah et al., 2019). Having a small set of objectives to improve can easily result in neglecting other dimensions of the issue and cause other problems while fixing the initial problem.

The questions related to the use of data science tools can also be placed in the broader context of migration ethics and political theory, where the responsibilities of the stakeholders, practical conditions and normative positions in the discourse are questioned (Carens, 2013; Owen, 2020). For the purposes of this chapter, we address a narrower set of concerns related to the use of specific technologies, but we acknowledge that the issues discussed here have a wide range of implications and the conceptual framing of these problems must be carefully considered. The last few years of increased technology use have taught us that some consequences of wide-range adoption of technologies are very difficult to predict. An example can be found in social media, where content filtering algorithms designed to improve user experience end up widening the gaps within the society and polarising it. Such risks are exacerbated when the computer scientists who design the algorithms are unaware of the nuances and debates around the topic and in related domains, treating

questionable assumptions as solid foundations. For this reason, a basic understanding of data ethics is crucial for the ethical practice of data science.

2.2.3 Ethics principles

In recent discussions in ethics of technology, data, and artificial intelligence (AI), ethics principles have been dominant. Since 2015, over a 100 sets of ethics principles have been published around the world just for AI practices.[1] This number is higher when we include data ethics principles and, more broadly, technology ethics principles. Each one of these sets cover a wide range of ethics principles such as transparency, privacy, diversity, and sustainability. While these sets of principles differ slightly from each other by putting emphasis on different principles, their basic structure is in line with the *principlism* approach developed in 1979 for research ethics with three *core principles*: respect for persons, beneficence, and justice.[2]

These core principles are loosely derived from theories in moral and political philosophy. Autonomy, minimisation of harm, and justice are argued to be intrinsically valuable within these theories.[3] They are valued for themselves and not as a means to achieve other goals—they are the goals of ethics. In comparison, *instrumental principles* such as privacy or transparency are valuable as a means to achieve these ends. For example, privacy allows us to exercise our autonomy without interference from others and protects us from harm; transparency allows us to acquire necessary understanding and information so that we can make rational and informed choices and thereby exercise our autonomy, and it helps upholding justice by uncovering unfair

[1] AI Ethics Lab, *Toolbox: Dynamics of AI Principles*, February 2020, see `https://aiethicslab.com/big-picture/`

[2] The Belmont Report (The National Commission for the Protection of Human Subjects of Biomedical and Behavioral Research, 1979) lays out the foundation of *principlism*, which is further developed in the seminal book, *Principles of Biomedical Ethics* (Beauchamp and Childress, 2013). While the Belmont Report lists three "basic ethical principles", principlism often lists four, dividing the principle of beneficence into beneficence and non-maleficence principles.

[3] In humanitarian aid and forced migration literature, "do no harm" principle makes up a core concept. This concept is derived from medical ethics. However, strictly speaking, "do no harm" is an ethically problematic principle. A good example is a life-saving surgery: Surgery necessarily involves harm (that is, cutting open the patient and risking their life) although it is justified by another risk of harm already present (that is, patient's medical condition which poses harm to the patient in the absence of surgery) and the potential of benefits which would outweigh the harm of surgery (that is, the cure). A strict "do no harm" principle would not allow a life-saving surgery, because while doing surgery involves doing harm, allowing the patient to die from their existing medical condition is not an active *doing* of harm. Clearly, this reasoning poses serious ethical problems. Therefore, "do no harm" functions as a loose title for a more nuanced ethical approach – that is, recognising and minimising the risk of harm due to (humanitarian or medical) intervention, and minimising the risk of overall harm while maximising overall benefits to ensure that benefits outweigh the harm. The exact procedures to achieve this end and the threshold of risk for justifiable interventions are further discussed in the relevant literature.

and unequal treatments. Such instrumental principles are not goals of ethics—that is, we do not value privacy or transparency simply for their own sake (Canca, 2020).

Why does this distinction of core and instrumental principles matter? Because in order to apply and utilise these principles, we need to understand which of them are interchangeable and which are fundamental. (We will return to this distinction in Section 2.4 with *the Box* tool.) Let's take the example of using aggregate anonymous mobile phone data to study the migration patterns in order to predict the spread of an infectious disease (Wesolowski et al., 2012). The goal here is to control the spread of the disease and thus minimise the harm. However, the data are not collected through individual consent, because collecting individual consent, even if possible, would slow down the process to such a degree that the project would fail. Is this an ethical violation? The answer depends on the purpose of consent and the other related ethical concerns. Consent, as an instrumental principle, serves to protect and promote individual autonomy by ensuring that individuals can make decisions regarding their life, their space, and their body. If the data are aggregated and anonymised properly, it cannot be re-identified or connected back to the individuals and therefore, its use will not directly affect the individual's life. In other words, through another instrumental principle—that is, privacy—the ethical need for consent to protect individual autonomy can be satisfied. One could still argue that if the project results in decisions that would negatively affect the individual's group (for example, travel bans for seasonal workers), the individual might have refused to allow their anonymised data to be used, if they were given the chance to do so beforehand.

At this point we are pitting core principles against each other. Will a certain action result in unfair discrimination of a group? Is the harm done to this group outweighed by the benefit this group and/or others might receive? Would protecting individual autonomy in this broad manner result in violation of the autonomy of others? In order to ensure that the project is ethically justifiable, these questions must be answered by a thorough ethics analysis that takes into account the empirical evidence and, if need be, appeals to the theories behind these principles to flesh out their specific demands. Once this analysis lays out the extent of ethical risks, we can determine the best course of action to mitigate these risks through project and algorithm design and safeguards for the use of the results.

2.3 Ethical issues in data science for migration and mobility

Here, we take a brief look at some of the most common ethical issues in data science and AI technologies. These include issues related to (1) consent, (2) de-identification, anonymisation, and re-identification, (3) black box, transparency, and explainability, (4) algorithmic bias, (5) dual use of technologies, and (6) complexity and risk assessment. As we sketch out the main ethical concerns with respect to each of these points, we also provide examples from the migration and mobility domain to illustrate the issues within the proper context of this book.

2.3.1 Consent

Consent is a practice to ensure individual autonomy is protected and promoted. Proper consent has three conditions: the individual consenting must be informed, rational, and the consent must be voluntary. The ethical and sometimes legal need for consent arises in various stages of data science and AI technologies. In research, consent plays a crucial role for the individual's participation and sharing of identifiable personal data.

In emerging technologies, a major problem with consent arises from a lack of understanding around the technology and its potential for harm and benefits. As we will discuss in detail in the next sections, difficulties in risk assessment of technologies, the lack of transparency in models, risks for re-identification of data, and dual use of technologies make it particularly difficult for individuals to engage in meaningful consent procedures. This is also due to the fact that even researchers are often unclear about these risks and benefits. Furthermore, power relations between authorities and vulnerable populations, as well as ownership of surveillance and data collection technologies affect practices of consent. For example, it is difficult for a refugee living in a camp to decline consent for an iris scan, if it is the only way to get food and fundamental help.[4]

Data collected and research conducted in online platforms through reliance on the terms and conditions of these platforms also raise a problem about consent. Since it has been shown that individuals cannot practically or reasonably read and understand all of the terms and conditions of all platforms they use, these agreements cannot constitute a proper consent (McDonald and Cranor, 2008).[5] From a regulatory perspective, these agreements are superseded by communication laws of the countries where these platforms reside, which may explicitly permit aggregated and anonymised processing of such data for humanitarian or research purposes.

Failure to properly obtain consent from refugees and asylum seekers may result not only in a violation of the right to privacy and self-determination, but also the right to liberty and security, and even the right to life. A striking example of this is the controversy about the United Nations Refugee Agency (UNHCR) collecting and sharing of Rohingya refugees' personal data with Bangladesh, which then shared it with Myanmar authorities. Allegedly, the UNHCR collected Rohingya refugees' personal data by having them sign a document, where a checkbox indicated that their data might be shared for repatriation purposes. However, this checkbox was in English only, and many refugees thought that they had to agree to the terms in order to get their identity cards. According to Human Rights Watch (HRW), this

[4] https://refugeesmigrants.un.org/ar/node/100042481

[5] https://www.pcmag.com/news/it-would-take-17-hours-to-read-the-terms-conditions-of-the-13-most-popular

constitutes a breach of the policies of the UNHCR, who subsequently released a statement to respond to the allegations.[6] [7]

2.3.2 De-identification, anonymisation, and re-identification

Consent is often not needed if an individual's personal data can be successfully de-identified or anonymised, meaning that the data cannot be associated with an individual anymore. However, this is easier said than done. As algorithms become more powerful and more datasets become available, it is increasingly difficult to completely anonymise the data and prevent re-identification (Sweeney, 2002; Waldo, 2016). Re-identification poses a problem to privacy and thereby to individual autonomy.

Anonymisation removes personal identifiers such as names, addresses, identifying numbers and keys from a database, but the remaining patterns may still be sufficient to identify a person uniquely. For example, behaviour can also be used as a biometric (de Montjoye et al., 2013). In some cases, researchers have ensured anonymity under certain assumptions, but later research developed methods invalidating these assumptions and made identification possible. One should be cautious of sharing an anonymous dataset with a party, who might possesses additional information that can be used to remove anonymity from (some of) the records. This does not mean that privacy is impossible. Aggregation of data is an additional step that can be taken to ensure privacy is secured. Appropriately anonymised and aggregated mobile phone data has been successfully used to estimate population distributions, where other data sources are scarce, outdated, or unreliable (Deville et al., 2014). For example, the Data for Development (D4D) Challenge opened a large set of mobile CDR to the research community with the aim of providing insights to policy makers about development (Blondel et al., 2012). The initial challenge contained communication graphs, which illustrated some of the properties of the social networks of the users. Later, Sharad and Danezis (2013) managed to de-anonymise part of these communication graphs by using graph-theoretic analyses. Subsequently, the following Data for Refugees (D4R) Challenge, which also made a large mobile CDR dataset from Turkey available to help creating insights about the Syrian refugees in Turkey (Salah et al., 2018), abstained from making the communication graph data available.

---

[6] https://www.hrw.org/news/2021/06/15/un-shared-rohingya-data-without-informed-consent

[7] https://www.unhcr.org/uk/news/press/2021/6/60c85a7b4/news-comment-statement-refugee-registration-data-collection-bangladesh.html

2.3.3 Black box, transparency, and explainability

The increasing use of complex machine learning models result in systems that are often referred to as "black boxes", where it is not clear how the system reaches its outcome. These systems lack transparency regarding how the model engages with the data to produce the outcome. In contrast to "black box" AI, *explainable* AI models are those whose outcome can be understood by humans. There are at least two different aspects to explainability that need considering. One is the explainability for the designer, who benefits from an improved understanding of how the model is reaching its decisions. The designer can then use these insights to identify potential issues. Such insights would allow the designer to reduce various risks of harm including those related to safety, security, and unjustified discrimination. The second aspect of explainability is related to the user, who thereby achieves a better understanding of the decisions of the system and can factor it in, in further deliberations. Enabling the user to engage with the system in this way promotes user autonomy.

Consider a (hypothetical) satellite imaging based AI system that predicts whether an area will experience climate related forced migration within the next ten years. The input of the AI system is a satellite image and the output is a binary decision about the target variable. Such a system may be created via supervised learning, where past data are used to predict future cases. A system that is more explainable for the designer could use the concept of attention and highlight those areas of the satellite image that were most important for its decision. The designer can then notice, for example, that the system invariably pays more attention to the centre of the image (which can be a bias in training image selection) and the designer can take measures to rectify this bias, for example by selecting randomly cropped images. On the other hand, a system that is more explainable for the user may point to the lack of forested areas and distance to water sources, which impacted the outcome. This can prompt the user to examine potential solutions and policies in dealing with the issue. In the first case, explainability is used to improve the AI system and protects it from incorrect and biased design and training, whereas the second case is about making the system more useful as a tool in decision making.

2.3.4 Algorithmic bias

Algorithmic bias occurs when the AI model systematically and unfairly discriminates against certain groups. This is different from the statistical concept of bias in machine learning, although related to it. To clarify this point, let us first explain the latter. Suppose we have a sample and we are estimating a parameter θ based on this sample. In machine learning terminology, the bias of our estimator is the difference between θ that we are trying to estimate and the expected value of the estimator (Alpaydin, 2020). For example, for a properly collected sample, the sample average is an unbiased estimator for the population mean, as the expected difference between them is zero. We want this kind of bias to be as small as possible. If we are estimating a population

parameter by means of a sample, a large bias can be created by incorrect sampling from the population. For example, if you would like to estimate the education level of a migrant group in a country, but sample only migrants from a small university town, you will create a bias; the expected average education level of your sample will be higher than the expected average education level of the population.

Algorithmic bias, on the other hand, is about systematic and unfair outcomes of an AI system. This might be intentional or (more often) unintentional. The bias can enter into the model through various routes: unrepresentative datasets, datasets that reflect existing social biases, discriminatory labelling of data, variables and proxies used within the models, and framing of the problem for the model. Some of these biases are easier to fix (such as unrepresentative datasets) and some of them are nearly impossible to eliminate (such as social biases within datasets). Algorithmic bias constitutes a serious problem for social justice because systems and technologies can amplify existing biases while hiding them behind a facade of mathematical objectivity. Algorithmic bias—and more generally, bias—is necessarily related to the concept of *justice* and *fairness*. Theories of justice in political philosophy offer us various definitions of justice and how to apply them. To achieve useful results, research on algorithmic bias and fairness in machine learning has to engage with this literature.

When investigating biases, a useful concept is *protected attribute*, which are data attributes that need to be tested for potential biases, such as ethnicity, gender, and age. Discrimination based on a protected attribute can happen in different ways. Mehrabi et al. (2019) describe five major categories, of which the most important are *direct discrimination*, which happens when a protected attribute directly leads to a biased outcome, and *indirect discrimination*, which happens when the system is not taking any protected attributes into account, but seemingly neutral attributes are acting as proxies for directed attributes. An example of the latter is the usage of postal codes, which may contain information about the ethnicity of subjects, if ethnic groups are more densely populated in certain areas of a city. Indirect discrimination requires extra attention, especially if the outcomes of data analysis are used for policy decisions.

Dealing with issues of bias requires looking at the data collection and analysis process critically and holistically. Who collects the data (*e.g.* a private vs. public entity) and whether this entity has a specific agenda, how the data are sampled (whether there are any structural issues and whether representativeness is ensured), and how the data are annotated, are all issues that can introduce bias even before the analysis is initiated.

Biases may also arise inadvertently from incorrect analysis. In data science, the collection of large scale data from a heterogeneous sample may lead to a phenomenon called *Simpson's paradox* (Blyth, 1972). Put simply, this issue happens when one has subgroups that exhibit biases that cancel each other out when the data are aggregated. Clemens (2020) provide an example of this for the emigration-income relationship, which shows very different patterns for the aggregate population compared to subgroups with different income levels. Another example related to demographics is given in Escalante et al. (2022), where a model was initially examined for gender

bias and found to have none, but further investigation into different age groups revealed a strong preference for younger women and older men in the system, which cancelled each other out in the age-aggregated analysis. Simpson's paradox can also happen with trend analysis, where aggregation can make trends appear or disappear (Alipourfard et al., 2018). Shuffling and randomisation based tests can be used to determine whether Simpson's paradox has any effect on the outcomes of a study (Lerman, 2018).

2.3.5 Dual use of technologies

Most technologies have a dual use: They can be utilised to benefit individuals and society but they can also be used against them. This issue of dual use of technology becomes bigger as it gets easier to modify the developed systems for ethically problematic purposes and have them adopted by actors with different political agendas. Dual use of technology is not a novel problem, but it is amplified due to the adaptability and the scale of deployment of emerging technologies.

Some of the most important ethical issues in the context of mobility arise from the dual use of technology. A population-level mobility tracking application could be the key to control a pandemic, but also a dangerous surveillance tool in the hands of an autocratic government that can use it to suppress and punish any actions against its authority (Oliver et al., 2020). Similarly, a remote sensing application designed to support and save refugee boats in the Mediterranean Sea can also be used to stop them before reaching European shores.[8]

It is worth noting that inadequate use of ethical instruments may end up masking the problems instead of helping to solve them. An ethical assessment committee that is available but not consulted, data management plans that are written in detail but not followed, and stakeholder consultations where function creep and dual use are hidden can function as ethics-washing, where ethics is used more as a lip-service than as a tool. To avoid circumventing ethical safeguards, "data protection by design and default" is a good practice to follow, where protection is incorporated in the design stage to prevent misuse. For example, in the Data for Refugees Challenge, which created a mobile CDR database from one million customers, including Syrian refugees and natives, this practice was followed to anonymise the data as a one-way transformation during data collection (Salah et al., 2018; Vinck et al., 2019). Since the original data, as well as the mapping, were removed before the database was shared, the anonymity of individuals would have been preserved even if a data breach were to compromise the dataset.

[8] See for example the discussions around humanitarian rescue operations in the Mediterranean reported by Médecins sans Frontières, *e.g.* `http://prez.ly/Yg0b`.

2.3.6 Complexity and risk assessment

A major source of difficulty in ethical design, development, and use of technologies comes from the complexity of human social dynamics and the difficulty of estimating technologies' effects on these dynamics. Many factors influence human behaviour, such as human mobility across the globe, and while it is possible to design experiments that control a large number of variables to study the effects of a few factors, the mutual interaction of these factors are difficult to model. The hallmark of complex systems is that a linear relationship between a set of causes and a set of effects is inadequate for modelling the system dynamics. Rather, these systems can be better conceptualised by stable and unstable attractors, equilibria, limit cycles, and bifurcations (Strogatz, 2018). When a technology is designed and put into use, it interacts with the society that uses it and changes it in unexpected ways. A good example is again social media, which ended up occupying a role in the society that no one could have predicted, as it slowly transformed practices of communication and commerce.

Complexity has three implications for ethical design in the domain of migration and mobility. The first is that the models and systems relying on big data analysis should be seen as potential agents of change. A technology that is initially designed to help in classifying asylum cases may create a benchmark that changes the behaviour of asylum seeking individuals. The second implication is that, like most complex systems, a control framework is required to properly monitor the system, where measurements should be obtained and continuously checked for drift. Finally, and most importantly, the conceptual tools of complex systems should be used effectively, instead of simpler but inadequately linear cause and effect explanations (Lauer, 2021).

In addition to the complexity problem, the problems of re-identification, black box, algorithmic bias, and dual use contribute to the difficulty of weighing the potential risks and benefits of systems and technologies. Difficulties in risk assessment pose a direct problem for minimisation of harm and they feed into other ethical issues such as consent and protection of vulnerable groups. When the risks are not well-understood and well-calibrated even by the researchers and developers, it becomes increasingly difficult for them to minimise the risk of harm to individuals, to vulnerable groups, and to the society as well as to explain these risks to individuals for their informed consent when sharing personal data or participating in research.

To give an example, the processing of mobile CDR in the Data for Refugees project clearly showed refugees working in areas in which there were no working permits issued to them (*e.g.* a large airport construction project). Publishing this result—as a purely scientific finding—at the time might have caused people to lose their jobs. Risk assessment must take into account the reality of vulnerable groups and not be restricted to an idealised situation. In doing so, it is also worth understanding what constitutes a vulnerable group. Traditionally, what constitutes a group in this context involves national, ethnic, religious, or racial ties. However, groups can be formed by internal or external perceptions of commonalities and identities (Verkuyten, 2018). Furthermore, as Kammourieh et al. (2017) point out, big data processing can create

groups by matching people via commonalities. Subsequently, group privacy also needs to be considered as a risk factor. Vulnerability can be legally defined (*e.g.* asylum seekers) but in many occasions, it can also be contextual and it needs to be considered in relation to power and authority. Political systems and agendas may sustain, create, or fail to alleviate vulnerabilities. Migrants and refugees may be disadvantaged in a certain context, they may have limited access to rights and services, even fundamental and human rights like the right for asylum. Big data and technology based projects need to be assessed from this perspective as well.

The complexity of technology-society interaction and the various facets of risk assessment also inform how we should understand and position ethics and ethics analysis with respect to technology development and use. Ethics analysis must take the gaps and uncertainties in risk assessment into account and evaluate technologies for their multiple impacts on individuals, groups, and the society. As technology transforms society and society transforms the technology, ethics analysis must function as a tool to detect and mitigate arising risks. To do that, ethics analysis must remain as a continuous and integral part of technology development and use.

2.4 Data ethics tools

This section serves as an entry point for useful tools and resources for evaluating a data science project from an ethics perspective. We strongly recommend working with ethics experts for any major initiative, as paying lip service to a few ethics guidelines will not be sufficient to thoroughly analyse and vet a project. Ethics expertise is necessary to engage with the ethics literature that would constitute the backbone of any ethics analysis and help utilise ethics tools fully.

In several major initiatives on mobility analysis, dedicated ethics committees were formed and operated at each stage of the project. For example in the Data for Development challenge project (Blondel et al., 2012), an ethics panel (DEEP - Data for Development external ethics panel) was created, which evaluated project proposals submitted to obtain a large mobile phone dataset, as well as the reports of the groups who did obtain the data, and finally wrote a report on the problematic issues and recommendations for such data (DEEP, 2015). This report, for example, acknowledged that "Big Data can enable understanding and modelling large scale human behaviour with a temporal and spatial granularity never achieved before," and pointed to several challenges:

- Local knowledge is necessary to interpret how people are using the technology in question. Forming local collaborations is an important aspect of the work.
- In low and middle income countries (LMIC), there is less awareness of risks of making personal data public, which does not mean that this is an opportunity to make more data public from LMIC, but that there is a need for taking extra precautions and protecting these people against risks they may not be aware of.
- Anonymisation should ensure a minimum security and trade-off granularity of detail in a risk-based approach.

- While most legislation is aimed at individual privacy, group privacy needs to be taken into account.
- Local legislation may be missing or insufficiently developed for digital data protection. International standards should be observed nonetheless.
- The statistical biases in proprietary data can be difficult to understand and quantify. Having reputable institutions analyse the data, and providing access to many scrutinising eyes can address this only partially.
- The publication of research results may touch on culturally sensitive subjects, or through some correct or incorrect inferences, may be harmful to the groups under study. The risk assessment should also consider the publications and public dissemination as potential risk elements.

As this report shows, having a dedicated ethics committee can help bring ethical problems and concerns to the surface, which would not necessarily have been acknowledged otherwise.

Data ethics tools and guidelines can bring order into the process of assessing ethical issues by emphasising various key concerns, and questioning the involvement and (potentially conflicting) aims of all stakeholders. Some such tools are *The Data Ethics Canvas* of Open Data Institute (ODI), *The Box* by AI Ethics Lab, and *Data Ethics Decision Aid* of Utrecht University.

The Data Ethics Canvas of Open Data Institute (ODI)[9] is a tool that groups a number of ethics-related questions in 15 headings, and prompts the researcher to answer each of these questions in turn. For example, one of the headings is "Negative effects on people", and under that heading, the following questions are asked:

- Who could be negatively affected by this project?
- Could the way that data are collected, used or shared cause harm or expose individuals to risk of being re-identified? Could it be used to target, profile or prejudice people, or unfairly restrict access (*e.g.* exclusive arrangements)?
- How are limitations and risks communicated to people? Consider: people who the data are about, people impacted by its use and organisations using the data.

The Box by AI Ethics Lab[10] is a tool for operationalising ethics principles (Fig. 2.1). It aims to help researchers, developers, and designers think through the ethical implications of the technologies that they are building. *The Box* is a simplified tool that lists important ethical concerns by putting 18 instrumental ethics principles in relation to three core principles: respect for autonomy, minimisation of harm and maximisation of benefits, and securing justice. For example, instrumental principles of human control, transparency, explainability, information, agency, consent, and privacy mainly help promote the core principle of respecting individual autonomy. Once we correctly distinguish between core and instrumental principles we can turn many vague AI principles into an operational checklist to guide practice, because the core principles reveal the underlying values practitioners should aim to achieve, while the instrumental principles offer various paths for achieving them. The categorisation

[9] `https://theodi.org/article/data-ethics-canvas/`

[10] `https://aiethicslab.com/the-box/`

of the instrumental principles in relation to specific core principles helps researchers and practitioners focus on different aspects of each core principle and offers a way to determine how to best satisfy the core principles by substituting or supporting one instrumental principle with another (Canca, 2020). In *The Box*, each of these instrumental principles are further detailed through prompt questions. Once the researcher engages with all of these questions, the tool also helps them visualise ethical strengths and weaknesses of the technologies that they are evaluating and enables visual comparison of these technologies.

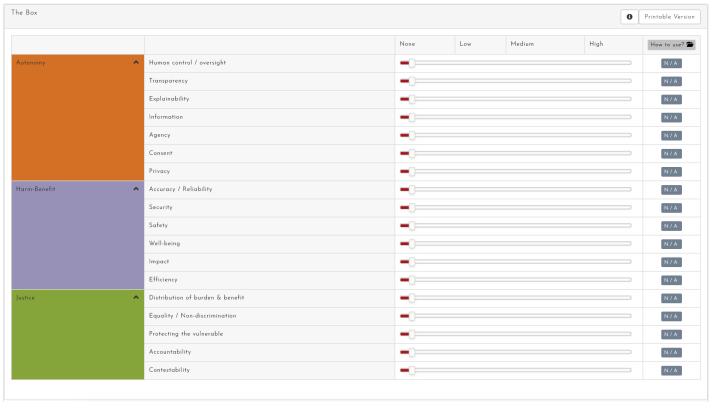

Fig. 2.1 The Box is a tool for visualising the strengths and weaknesses of a technology from an ethical perspective.

Another example is the *Data Ethics Decision Aid* (DEDA), developed at Utrecht University for reviewing public projects with social impact using large scale citizen data (Franzke et al., 2021). Especially for municipalities and local governance, data-driven management is an important tool, because near real-time monitoring helps with rapidly responding to the needs of the city. The authors point out the fact that legal frameworks and regulations are inadequate to deal with all the issues related to such data usage, and that there are legal usage instances which are ethically problematic. The main difference of DEDA is that it defines a number of roles (such as project lead and policy officer) within the organisation, and associates specific actions to these roles. DEDA also structures its activities around asking a pre-determined set of questions, which are organised into the headings of data related considerations (collection algorithms, source, data use, including anonymisation and visualisation, data storage, including access, sharing, reusing and re-purposing) and general considerations (responsibility, communication, transparency, privacy, and bias).

2.5 Legal risks for data science in migration and mobility

Ethical guidelines are necessary, because legal frameworks do not cover everything related to the development and use of technological solutions. Legal frameworks, even the human rights framework, leave many questions open in real-life practice (Canca, 2019). There are many cases where data collection and processing practices might be legal, but not ethical, such as using excessive but legally permitted surveillance measures to observe the behaviour of employees (Franzke et al., 2021). It is therefore important to separate the ethical and the legal issues. We discuss the latter in this section.

Legally, researchers and developers only make themselves liable if they fail to comply with the national and international standards and regulations they are subject to. While in some cases they must deal with issues of dual loyalty (*e.g.* involvement of multiple jurisdictions), they may expect to be exempt from any individual legal responsibility for the possible damages their products cause if they are operated diligently. However, if such damages and human rights violations are the result of governmental operations of the systems the developers created, they can still result in state responsibility under international human rights law, or administrative liability under the relevant national jurisdiction.

We discuss legal aspects in three stages, namely, data collection, algorithm development, and actual deployment and policy. For data collection, the most important concepts are "consent," "data protection," and "identity". For algorithm development, we will discuss "fairness," "transparency," and "bias," and for deployment, "accountability" and "power". We do not attempt an exhaustive treatment of these complex issues here, we merely point to some of the important concepts and considerations.

2.5.1 Data collection, processing, and sharing

The central issue in big data, from a legal perspective, is "consent," without which data collection (and processing) may be illegal. In Europe and in the US, the legal traditions lead to different paths when it comes to consent (Boehm, 2015). The European Union (EU) and the Council of Europe (CoE) regulations on the issue, such as General Data Protection Regulation (GDPR)[11] asserts that *data protection* and the Data Protection Convention, emphasise substantive law guarantees, including fairness, lawfulness, adequacy, and purpose limitation. The latter implies that collected data should not be further processed in a way incompatible with the purpose for which it was collected. The US regulations, however, lack such guarantees. In addition, the US and Canadian regulations treat foreigners outside of their territory differently by not applying the legal standards applicable in their

[11] `https://gdpr.eu/article-4-definitions/`

countries (Hayes, 2017). As a result, what is considered as a legal consent in a given context may differ greatly according to jurisdiction.

The second important issue is "data protection," which pertains to the security of the collected and stored data. The GDPR asserts that *data protection* means keeping data safe from unauthorised access and *data privacy* means empowering users to make their own decisions about who can process their data and for what purpose. The right to privacy is subsequently also connected to the user's understanding of how their private data are shared with other parties, and to consent to any such sharing of data.

Especially in cases where data from vulnerable populations are collected, data protection has serious legal implications. Refugees, human trafficking victims, political asylum seekers are examples of these cases. If a data leak may jeopardise or endanger the data owners, regulations like GDPR require that the data are kept encrypted and secure, and prevent data from being shared with third parties that do not meet the required standards. Failure to comply with these standards may result in legal responsibilities for states or individuals. Protection extends beyond the lifetime of a project, and while designing the project, one should regulate whether data will be archived or destroyed at the end of it. In many cases, there are legal mandates for the maximum storage duration of particular types of data.

The third issue is "identity," which from a legal perspective, is connected to the processing of personal data, as the latter is defined as any information that relates to an identified or identifiable living individual. From an EU perspective, Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 (GDPR) define a set of human rights related principles and rules regarding the processing of personal data. As regulated under Article 5 of the GDPR, these principles include lawfulness, fairness and transparency, purpose limitation, data minimisation, accuracy, storage limitation, integrity, confidentiality, and accountability. Additionally, any processing of personal data must be based on the data subject's consent (in cases of special categories data, this consent must be, as a rule, specific), and be subject to the principle of proportionality. While GDPR is an EU regulation, these principles have worldwide ramifications, particularly because personal data can only be transferred from the EU to third countries in compliance with the conditions and standards set out in GDPR.

On 14 May 2019, the EU adopted Regulations 2019/817[12] and 2019/818[13], establishing a framework for the interoperability between EU information systems in the Area of Freedom, Security, and Justice. These regulations aim to establish an interoperability regime between various existing EU databases by creating four new components, which will be accessible to border control and law enforcement authorities of EU member states and to Europol and Interpol. Blasi Casagran (2021) notes that these components actually constitute new databases since they will be processing and storing data in a structured manner, and incorporate new objectives.

[12] https://eur-lex.europa.eu/eli/reg/2019/817/oj

[13] https://eur-lex.europa.eu/eli/reg/2019/818/oj

As such, they may constitute a separate interference with human rights, particularly the right to privacy of the subjects.

The interoperability regulations also raise concerns regarding the right to non-discrimination of third country nationals that travel or migrate to EU countries. The merger of data affected by the components created through these regulations only concerns third country nationals, and might constitute, in and of itself, a discriminatory attitude toward them by subjecting them to other standards than EU citizens, and by regarding them as potential criminals (Blasi Casagran, 2021).

2.5.2 Algorithm development

Personal data regarding migrants, refugees, and asylum seekers may be subject to automated decision making processes, some of which may be using AI algorithms. The algorithms intended to be used in migration, asylum, or border-control management may have an adverse impact on the fundamental rights of the subject. We have already discussed algorithmic fairness and transparency earlier in this chapter. Here, we briefly mention some additional points.

A proposal for an EU Regulation laying down harmonised rules on AI addresses such algorithms as "high-risk" AI systems and aims to establish requirements and standards for the development, commercialisation, and use of those AI systems "that pose significant risks to the health and safety or fundamental rights of persons".[14] The proposal stresses the importance of accuracy, transparency, and the non-discriminatory nature of these systems when intended to be used as polygraphs and similar tools or to detect the emotional state of a person, for risk assessment of persons entering the territory of an EU member state or applying for visa or asylum, for verifying the authenticity of their documents, or for determining the eligibility of their applications. If put into force, developers and users of such high-risk AI systems will have to meet certain requirements, such as establishing a risk management system, guaranteeing high-quality data, transparency, human oversight, accuracy, robustness, and cybersecurity.

Developers of high-risk AI systems should be particularly aware that AI systems may be inadequate in dealing with non-standard situations. In such cases, the lack of an effective human oversight may lead to an infringement of fundamental rights. In order to reduce or eliminate legal risks arising from conformity requirements for high-risk systems, the system should incorporate appropriate human-machine interface tools to allow effective human oversight. The human overseer must be able to correctly interpret and override the output of the system and to refrain from overly relying on it. To this end, developers and producers should provide necessary information, tools, and education for users.

[14] https://digital-strategy.ec.europa.eu/en/library/proposal-regulation-laying-down-harmonised-rules-artificial-intelligence-artificial-intelligence

2.5.3 Deployment and policy

Legal problems related to the deployment of big data systems are mainly associated with powerful users, which may be government agencies or private corporations. Algorithmic approaches, based on big data and AI technologies, are used to inform policies of governments and intergovernmental organisations. In the area of migration and mobility, an example is IOM's Displacement Tracking Matrix (DTM), which relies on several big data sources, like mobile phone data and social media analysis to monitor movements of people across the globe. According to its website, DTM "gathers and analyses data to disseminate critical multi layered information on the mobility, vulnerabilities, and needs of displaced and mobile populations that enables decision makers and responders to provide these populations with better context specific assistance."[15]

Two important considerations here are accountability, and relations of power between data collectors and data owners. In this context, accountability (and oversight) are about the existence of regulatory mechanisms to protect people against unfair decisions taken by automated systems. When formulated as a supervised learning problem, an automated decision making system may be designed for example to make a recommendation about the admittance of a person through a border. The mathematical formulation of the system and the preparation of its parameters (i.e. training) requires a cost function, where each error the system makes has a certain cost. Admitting someone incorrectly may not have the same cost as denying someone entry incorrectly. One could argue that the cost of an error is a human cost, and it is perhaps impossible to reduce that to a numeric value. Accountability is about holding the organisations operating such systems accountable for the errors.

While the processing of personal data may be used to better allocate migration management services and to protect potential victims of human trafficking (Beduschi, 2017), it may also cause human rights infractions. Particularly in the case of processing methods utilising AI, automated decisions may impact migrants, refugees, and asylum seekers disproportionately. It should be considered that such individuals may be disadvantaged regarding their ability to recognise that they are dealing with an AI, to be informed about their rights to submit the decision to a human supervisor, and to use such rights in an effective way.

The already asymmetrical power structure between the public authorities and the migrants can be furthered through the use of advanced technologies. Countries receiving large numbers of migrants use migration management approaches that are increasingly relying on technology (Geiger and Pécoud, 2010). Molnar (2019*b*) argues that such technologies do not prioritise human rights and that states deliberately keep international regulation to a minimum. Risk assessment techniques utilising big data for immigration and border control may result in criminalising migration and putting vulnerable groups in jeopardy (Beduschi, 2017). Additionally, accountability gaps are created when states outsource part of their responsibilities in this area to

[15] `www.globaldtm.info`

private companies. Large scale surveillance data are being collected via satellites and drones for border control (*e.g.* Frontex and the European Border Surveillance System in the EU), leading to (further) militarisation of border control agencies and to the entrenchment of the image of migrants as a threat to be averted (Csernatoni, 2018).

The use of technology for the detection and pushback of migrants, refugees, and asylum seekers has ramifications regarding the extraterritorial jurisdiction of the state utilising such methods. It is possible for states to employ a humanitarian technology to detect and track asylum seekers, and ultimately to implement pushback strategies. In *Hirsi Jamaa and others v. Italy*, the European Court of Human Rights decided that migrants intercepted by Italian forces in open sea and then deported to Libya fell under the jurisdiction of Italy, although they never entered Italian territorial waters. Furthermore, the Court decided that the deportation to Libya constituted a violation of article 3 of the European Convention of Human Rights due to the risk of suffering ill-treatment in that country.[16] Marin and Krajčíková (2016) argue that this judgement had an effect on European countries increasingly seeking the cooperation of North African countries. Nonetheless, the use of data processing technologies to pushback vulnerable migrants, particularly refugees and asylum seekers may result in a breach of the *non-refoulement* principle if the people in question would be subject to torture, ill-treatment, political persecution, or death penalty in the country they would be deported to. Additionally, such actions could violate positive obligations of states to protect the life and liberty of persons arising from international law and human rights law (Molnar, 2019*b*).

One of the risks associated with the processing of big data is that connected databases create a number of cross-references, which, when used out of context, may result in discriminatory or restrictive decisions by law enforcement or border management authorities. The EU interoperability regulations mentioned above constitute a good example for this. Any official accessing one of the components would be pinged with a match if the data in question is included in any of the databases, without needing to access the contents of that database. In such cases, even a faulty or legally insignificant match (i.e. cases of double identity without any illegal reason), may result in a denial of entry to a country, if the decision-making official pursues a risk-averse attitude. Researchers should bear in mind that legal standards should serve to protect the individual from abuse or misuse of power by the authorities, and allow for safeguards to prevent such consequences.

2.5.4 Data ethics and the Global Compact for Safe, Orderly, and Regular Migration

The relevance of data ethics in the context of migration and refugees is emphasised also in the Global Compact for Safe, Orderly and Regular Migration (GCM), as policies at the United Nations (UN) level have the potential to affect the legal system

[16] `http://hudoc.echr.coe.int/spa?i=001-109231`

of many countries. On 19 December 2018, the UN General Assembly adopted a GCM.[17] The GCM comprises 23 objectives and commitments based on the 2016 New York Declaration for Refugees and Migrants. Objective 1 of the GCM addresses the need to collect and utilise accurate and disaggregated data as a basis for evidence-based policies. Commitments following this objective call for a "comprehensive strategy for improving migration data at the local, national, regional and global levels" through "harmonising data collection methodologies" and "strengthening analysis and dissemination of migration-related data and indicators". The same objective also supports "further development of and collaboration between existing global and regional databases and depositories".

Additionally, the GCM addresses the needs for an integrated, secure, and coordinated border management policy (Objective 11) and for strengthening "certainty and predictability in migration procedures for appropriate screening, assessment and referral" (Objective 12), among others.

While these objectives and commitments allow and encourage the collection, processing, and sharing of migration data to prevent illegal migration, the GCM also emphasises the need to protect the human rights and the principles of non-discrimination and non-regression for all migrants. Particularly Objective 17 includes a commitment to "establish mechanisms to prevent, detect and respond to racial, ethnic and religious profiling of migrants by public authorities, as well as systematic instances of intolerance, xenophobia, racism, and all other multiple and intersecting forms of discrimination, in partnership with national human rights institutions, including by tracking and publishing trend analyses, and ensuring access to effective complaint and redress mechanisms".

As a legally non-binding instrument, the GCM cannot set legal obligations to states or individuals, but may be considered as a benchmark for future policies of states willing to implement it. Already, a majority of UN Member States have made it a priority to eliminate discriminatory procedures against migrants and to introduce a concept of responsibility for states that follow this commitment (Guild, 2018).

2.6 In conclusion

Data science in migration and human mobility involves dealing with a number of ethical questions and concerns in every stage of a project. These ethical questions cannot be separated from the development and design of the project, since they arise in defining the goal of the project, in data collection, labelling, and storage, in interactions with research subjects, and in the application of the results in real-life. For that reason, ethics must be considered as an integral part of data science projects from the very start as something that requires not only care but also problem-solving skills using ethics tools.

[17] `https://www.iom.int/global-compact-migration`

Ethical and legal concerns regarding human mobility and migration are deeply intertwined since new legal regulations arise from developing emerging practices, which in turn take existing ethical standards into consideration. Discussions for determining which regulations to put in place must first necessarily engage in ethical debate on what would be the "right" and "fair" practices in development and use of technologies. Regardless how detailed the legal framework becomes, there are and will always be questions that fall into grey areas and require ethical decision-making during the technology development and deployment process. National and international legal instruments do not always and immediately result in a binding law, but in many cases involve ethical self-assessment procedures or general policy guidelines. As such, it is imperative for researchers and developers to be aware of the ethical risks associated with their analyses and products, particularly if the systems are expected to interfere with human rights.

2.7 Document repositories and further reading

There are several documents that are relevant and useful when it comes to addressing legal and ethical aspects of big data projects. User license agreements are necessary for datasets collected or shared by private companies. GovLab maintains a repository of agreements to form a starting point for any new project, which can tremendously lighten the load of the legal team.[18] The DARIAH project collects a repository of consent forms for users.[19] During the course of a project, it is important that data access is clearly regulated, which is typically achieved with a Data Management Plan (DMP). Many universities and funding agencies have their own standards or templates for these documents, and Stanford University Library has a useful repository of DMPs.[20]

Applied ethics must take into account the context within which the ethical questions arise. We have mentioned several ethics frameworks for reviewing the projects. In order to use these tools properly, we need to engage with the demands and details of the context. A large amount of resources are available on the websites of Ethics in Context,[21] the Oxford Institute of Internet,[22] and Stanford Institute for Human-Centered AI.[23]

The concept of privacy, in its multiple forms, is extensively discussed in Solove (2006), and neatly summarised in Kitchin (2016). Gambs et al. (2014) is a good

[18] https://www.thegovlab.org/project/project-contracts-for-data-collaboration

[19] https://www.dariah.eu/2020/09/17/dariah-eldah-consent-form-wizard/

[20] https://library.stanford.edu/research/data-management-services/data-management-plans

[21] https://c4ejournal.net/

[22] https://www.oii.ox.ac.uk/

[23] https://hai.stanford.edu/

starting point for understanding risks of sharing geo-located data, and how these can be de-anonymised.

The domain specific terminology and knowledge are very important; for data scientists not familiar with the area of migration and mobility, UNHCR and IOM provide white papers and reports that can serve as a starting point. A key document in assessing proportionality and power usage for issues with regard to vulnerable populations is the "Principles and Guidelines, supported by practical guidance, on the human rights protection of migrants in vulnerable situations," prepared by United Nations Human Rights Office of the High Commissioner.[24] The 1951 Geneva Convention and its 1967 Protocol are the most important legal documents to clarify definitions and rights for refugees, and to determine the obligations of the states which signed the protocol.[25] However, there are special cases to consider. For example, while Turkey is party to the 1951 Geneva Refugee Convention, it only acknowledged "refugee" status for people originating from Europe. Consequently, the Syrian refugees in Turkey were officially considered "temporarily protected foreign individuals".

Several researchers published seminal work on the role of AI and big data technologies in migration management. Ana Beduschi's work discusses both the role of technology (Beduschi, 2017, 2021), and the concept of digital identity (Beduschi, 2019). Molnar and Gill (2018) wrote an influential report on automated decision making employed in Canada's immigration and refugee system, and Akhmetova (2020) extends this discussion with the concept of an 'invisible border wall'. Other work by Molnar contributes important ideas from a human rights perspective (Molnar, 2019*a*,*b*).

On the data science side, proper statistical or machine learning based experiment design is a key issue. Mehrabi et al. (2019) provide a comprehensive survey of bias and fairness in machine learning and algorithmic decision making, and Lepri et al. (2018) provide an overview of technical solutions to enhance fairness, accountability, and transparency in such settings.

Acknowledgments

This study is supported by European Union's Horizon 2020 Research and Innovation Programme under grant agreement No 870661.

[24] https://www.ohchr.org/Documents/Issues/Migration/PrinciplesAndGuidelines.pdf

[25] https://www.unhcr.org/1951-refugee-convention.html

References

Akhmetova, R. (2020), 'Efficient discrimination: On how governments use artificial intelligence in the immigration sphere to create and fortify 'invisible border walls'', *Centre on Migration, Policy and Society Working Paper* **149**.

Alipourfard, N., Fennell, P. G. and Lerman, K. (2018), Can you trust the trend? discovering Simpson's paradoxes in social data, *in* 'Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining', pp. 19–27.

Alpaydin, E. (2020), *Introduction to machine learning*, 4th edn, MIT press.

Beauchamp, T. L. and Childress, J. F. (2013), *Principles of Biomedical Ethics*, 7th edn, Oxford University Press.

Beduschi, A. (2017), 'The big data of international migration: Opportunities and challenges for states under international human rights law', *Georgetown Journal of International Law* **49**, 981–1018.

Beduschi, A. (2019), 'Digital identity: Contemporary challenges for data protection, privacy and non-discrimination rights', *Big Data & Society* **6**(2), 2053951719855091.

Beduschi, A. (2021), 'International migration management in the age of artificial intelligence', *Migration Studies* **9**(3), 576–596.

Blasi Casagran, C. (2021), 'Fundamental rights implications of interconnecting migration and policing databases in the EU', *Human Rights Law Review* **21**(2), 433–457.

Blondel, V. D., Esch, M., Chan, C., Clérot, F., Deville, P., Huens, E., Morlot, F., Smoreda, Z. and Ziemlicki, C. (2012), 'Data for development: the D4D challenge on mobile phone data', *arXiv preprint arXiv:1210.0137*.

Blyth, C. R. (1972), 'On Simpson's paradox and the sure-thing principle', *Journal of the American Statistical Association* **67**(338), 364–366.

Boehm, F. (2015), 'A comparison between US and EU data protection legislation for law enforcement purposes', *European Parliament, Study for the LIBE Committee*.

Canca, C. (2019), 'Human rights and AI ethics – why ethics cannot be replaced by the UDHR', *United Nations University, Centre for Policy Research*.

Canca, C. (2020), 'Operationalizing AI ethics principles', *Communications of the ACM* **63**(12), 18–21.

Carens, J. (2013), *The ethics of immigration*, Oxford University Press.

Clemens, M. A. (2020), 'Migration from developing countries: Selection, income elasticity and Simpson's paradox', *Centro Studi Luca d'Agliano Development Studies Working Paper* **465**.

Csernatoni, R. (2018), 'Constructing the EU's high-tech borders: FRONTEX and dual-use drones for border management', *European Security* **27**(2), 175–200.

De Montjoye, Y.-A., Hidalgo, C. A., Verleysen, M. and Blondel, V. D. (2013), 'Unique in the crowd: The privacy bounds of human mobility', *Scientific reports* **3**(1), 1–5.

DEEP (2015), Data for development Senegal: Report of the external review panel, Technical report, Institute of Business Ethics.

Deville, P., Linard, C., Martin, S., Gilbert, M., Stevens, F. R., Gaughan, A. E., Blondel, V. D. and Tatem, A. J. (2014), 'Dynamic population mapping using mobile phone data', *Proceedings of the National Academy of Sciences* **111**(45), 15888–15893.

Escalante, H. J., Kaya, H., Salah, A. A., Escalera, S., Güçlütürk, Y., Güçlü, U., Baró, X., Guyon, I., Jacques, J. C., Madadi, M. et al. (2022), 'Modeling, recognizing, and explaining apparent personality from videos', *IEEE Transactions on Affective Computing*.

Floridi, L. and Taddeo, M. (2016), 'What is data ethics?', *Phil. Trans. R. Soc. A* **374**, 20160360.

Franzke, A. S., Muis, I. and Schäfer, M. T. (2021), 'Data ethics decision aid (DEDA): a dialogical framework for ethical inquiry of AI and data projects in the Netherlands', *Ethics and Information Technology* pp. 1–17.

Gambs, S., Killijian, M.-O. and del Prado Cortez, M. N. (2014), 'De-anonymization attack on geolocated data', *Journal of Computer and System Sciences* **80**(8), 1597–1614.

Geiger, M. and Pécoud, A. (2010), The politics of international migration management, *in* M. Geiger and A. Pécoud, eds, 'The politics of international migration management', Springer, pp. 1–20.

Guild, E. (2018), 'The UN global compact for safe, orderly and regular migration: What place for human rights?', *International Journal of Refugee Law* **30**(4), 661–663.

Hayes, B. (2017), 'Migration and data protection: Doing no harm in an age of mass displacement, mass surveillance and "big data"', *International Review of the Red Cross* **99**(904), 179–209.

Haznedar, B., Peyton, J. K. and Young-Scholten, M. (2018), 'Teaching adult migrants', *Critical Multilingualism Studies* **6**(1), 155–183.

Kammourieh, L., Baar, T., Berens, J., Letouzé, E., Manske, J., Palmer, J., Sangokoya, D. and Vinck, P. (2017), Group privacy in the age of big data, *in* 'Group Privacy', Springer, pp. 37–66.

Kitchin, R. (2016), 'The ethics of smart cities and urban science', *Philosophical transactions of the royal society A: Mathematical, physical and engineering sciences* **374**(2083), 20160115.

Lauer, D. (2021), 'You cannot have AI ethics without ethics', *AI and Ethics* **1**(1), 21–25.

Lepri, B., Oliver, N., Letouzé, E., Pentland, A. and Vinck, P. (2018), 'Fair, transparent, and accountable algorithmic decision-making processes', *Philosophy & Technology* **31**(4), 611–627.

Lerman, K. (2018), 'Computational social scientist beware: Simpson's paradox in behavioral data', *Journal of Computational Social Science* **1**(1), 49–58.

Mamei, M., Cilasun, S. M., Lippi, M., Pancotto, F. and Tümen, S. (2019), Improve education opportunities for better integration of Syrian refugees in Turkey, *in* 'Guide to Mobile Data Analytics in Refugee Scenarios', Springer, pp. 381–402.

Marin, L. and Krajčíková, K. (2016), Deploying drones in policing southern european borders: constraints and challenges for data protection and human rights, *in* 'Drones and unmanned aerial systems', Springer, pp. 101–127.

McDonald, A. M. and Cranor, L. F. (2008), 'The cost of reading privacy policies', *A Journal of Law and Policy for the Information Society* **4**(3), 543–568.

Mehrabi, N., Morstatter, F., Saxena, N., Lerman, K. and Galstyan, A. (2019), 'A survey on bias and fairness in machine learning', *arXiv preprint arXiv:1908.09635*.

Molnar, P. (2019*a*), 'New technologies in migration: human rights impacts', *Forced Migration Review* **61**, 7–9.

Molnar, P. (2019*b*), 'Technology on the margins: AI and global migration management from a human rights perspective', *Cambridge International Law Journal* **8**(2), 305–330.

Molnar, P. and Gill, L. (2018), 'Bots at the gate: a human rights analysis of automated decision-making in Canada's immigration and refugee system', *Citizen Lab and International Human Rights Program*.

Oliver, N., Colizza, V., de Cordes, N., Fraiberger, S., Koebe, T., Lehmann, S., Murillo, J., Pentland, A., Pham, P., Pivetta, F. et al. (2020), 'Mobile phone data for informing public health actions across the COVID-19 pandemic life cycle', *Science Advances* **6**(23).

Owen, D. (2020), *What do we owe to refugees?*, Polity.

Salah, A. A., Altuncu, M. T., Balcisoy, S., Frydenlund, E., Mamei, M., Akyol, M. A., Arslanlı, K. Y., Bensason, I., Boshuijzen-van Burken, C., Bosetti, P. et al. (2019), Policy implications of the D4R challenge, *in* A. A. Salah, A. Pentland, B. Lepri and E. Letouzé, eds, 'Guide to Mobile Data Analytics in Refugee Scenarios', Springer, pp. 477–495.

Salah, A. A., Pentland, A., Lepri, B., Letouzé, E., Vinck, P., de Montjoye, Y.-A., Dong, X. and Dağdelen, Ö. (2018), 'Data for refugees: The D4R challenge on mobility of Syrian refugees in Turkey', *arXiv preprint arXiv:1807.00523*.

Sharad, K. and Danezis, G. (2013), De-anonymizing D4D datasets, *in* 'Workshop on hot topics in privacy enhancing technologies'.

Solove, D. (2006), 'A taxonomy of privacy', *University of Pennsylvania Law Review* **154**(3), 477–564.

Strogatz, S. H. (2018), *Nonlinear dynamics and chaos with student solutions manual: With applications to physics, biology, chemistry, and engineering*, CRC press.

Sweeney, L. (2002), 'Achieving k-anonymity privacy protection using generalization and suppression', *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems* **10**(05), 571–588.

The National Commission for the Protection of Human Subjects of Biomedical and Behavioral Research (1979), 'The Belmont report: Ethical principles and guidelines for the protection of human subjects of research'.

Verkuyten, M. (2018), *The social psychology of ethnic identity*, Routledge.

Vinck, P., Pham, P. N. and Salah, A. A. (2019), "Do No Harm" in the age of big data: Data, ethics, and the refugees, *in* A. A. Salah, A. Pentland, B. Lepri and E. Letouzé, eds, 'Guide to Mobile Data Analytics in Refugee Scenarios', Springer, pp. 87–99.

Waldo, J. (2016), Big data and the social sciences: Can accuracy and privacy co-exist?, *in* 'Data for Policy 2016 - Frontiers of Data Science for Government: Ideas, Practices and Projections (Data for Policy)'.

Wesolowski, A., Eagle, N., Tatem, A. J., Smith, D. L., Noor, A. M., Snow, R. W. and Buckee, C. O. (2012), 'Quantifying the impact of human mobility on malaria', *Science* **338**(6104), 267–270.