

Democratising the digital revolution: the role of data governance

Sylvie Delacroix^{1 2}, Joelle Pineau³, and Jessica Montgomery¹

¹ University of Birmingham, Birmingham, B15 2TT, UK

² The Alan Turing Institute, London, W1 2DB, UK

³ McGill University, Quebec, H3A0G4, Canada

1 Introduction

Data is at the heart of today's AI. The machine learning techniques enabling many of the field's most recent advances and impressive applications leverage large amounts of data to extract insights that form the basis of new products or services. These systems have the potential to support hugely beneficial societal outcomes across a range of spheres of life – from improving healthcare services,¹ to increasing access to transport,² to helping tackle the major challenges posed by climate change.³ As AI technologies advance at a rapid pace, action is needed today to develop and implement governance structures to ensure that the benefits of AI are shared across society.

This chapter explores more specifically the role that data governance can play in shaping the development of AI technologies. It starts by considering how the role of law and governance systems in the digital environment is shifting, prompted by investigations or public incidents that have exposed the negative or unintended consequences of data use for both individuals and society. As the 'wild west' view of the digital sphere as an ungoverned, or ungovernable space, becomes increasingly outmoded, the chapter considers how policymakers and legislators are increasingly seeking means through which to assert social values in digital systems.

With a variety of legal and policy structures already seeking to influence patterns of data use and technology development, this chapter then briefly reviews recent legislative and policy activities, noting that – despite recent efforts – gaps in the policy landscape remain. Finding that new forms of bottom-up data sharing arrangement are needed to enhance democratic governance of data use, the chapter concludes by exploring the role of data trusts as a vehicle for leveraging the power associated with data aggregation.

2 Data for intelligence: the role of data governance in creating AI that benefits humanity

While the term AI for many conjures images of human-like intelligence, the type of intelligence that comes from today's techniques is different. The combination of advanced statistics and computing power that underpins many of the most successful AI technologies is perhaps more analogous to the human immune system than it is to human cognition.⁴ By processing data, these technologies are able to detect signals in the environment, which are not otherwise easily identifiable, and generate automatic responses to well-defined (and typically narrowly scoped) prediction tasks.

With data central to the development of AI, data governance will need to be central to any system seeking to encourage its trustworthy development and deployment. Effective data governance plays a role in both unlocking the value of data – enabling individuals and organisations to share data to support economic and social wellbeing – and protecting individuals, communities and society from the vulnerabilities that can be associated with the use of data, in particular the use of sensitive personal data. These vulnerabilities can relate to the privacy of such data, the development of data-enabled systems that reinforce discrimination on the basis of personal characteristics, or the potential for digital systems to shape the choices made by an individual in ways that undermine their agency both on- and off-line.

As the digital economy grows, and as data-enabled products and services become embedded in many daily activities, policymakers are grappling with questions about how best to manage such vulnerabilities. After an 'annus horribilis' for AI, in which a range of news stories laid bare the ways in which these new uses of data can leave individuals or groups exposed to harm,⁵ governments are increasingly looking for innovative governance mechanisms. The aim is to find ways of unlocking responsible data sharing while embedding legal and ethical practices that reduce the risk of harm and protect individual rights and freedoms.

There already exist legal instruments that seek to protect individual rights. Taking its roots in a human rights framework, the European Union's General Data Protection Regulation, for example, sits alongside a range of other legal instruments aimed at managing intellectual property, preserving copyright, and protecting privacy. Together, these create a constellation of individual rights and protections, and define circumstances and means through which individuals can assert those rights.

The nature of today's digital environment puts pressure on these existing systems. Designed for decisions of significant personal or social impact, these legal frameworks are not as well-equipped to handle the collective aspects of data sharing and manage the vulnerabilities that arise from the cumulative ways in which individuals share their data. Given that multiple algorithmic systems often act in parallel -- each leveraging parcels of data that inform seemingly insignificant decisions which become collectively significant -- tackling these vulnerabilities demands sophisticated

measures to *anticipate* the many risks of data use or potential failures in governance systems (in contrast to the current post-hoc, harm-remedying approach).

The latter, top-down approaches to constraining the use of data cannot by themselves create the conditions that support the beneficial use of data and AI. Today's challenge is therefore to bridge the gap between society's data sharing aspirations on the one hand and rights-protecting concerns on the other. This challenge creates a demand for new tools that can limit – or redistribute – technological and economic power. Inspired by discussions between Paul Nemitz (European Commission), Neil Lawrence (University of Cambridge), Nigel Shadbolt (Open Data Institute) and Lise Getoor (University of California at Santa Cruz), this chapter considers the infrastructures that could contribute to this democratisation of data governance.

3 The role of law and governance in the digital environment

3.1 Understanding the lessons from recent history

The last ten years have seen a rapid proliferation of data uses, and the growth of a vibrant global digital economy. Though the benefits of data use can be difficult to quantify, research suggests that on average, the use of data analytics improves company performance resulting in 5-6% higher output and productivity.⁶ Personal data has been a source of value in this economy. As more data is collected about individuals from a wider variety of sources – from online shopping, social media, fitness tracking devices, or mobile phone apps – it is increasingly possible for companies to develop a rich picture of daily life from the data trail left by each individual. This granular data is in turn relied on to build personal profiles whose predictive power is easily monetised in today's economy.

While bringing many benefits, these uses of data are exposing vulnerabilities. These accrue:

- To individuals, with examples of sensitive data about an individual's personal characteristics being inferred from seemingly innocuous information, as datasets are analysed in new ways, creating risks that individuals might inadvertently disclose private information;⁷
- To groups, as the social inequalities embedded in datasets at the point of collection are reinforced in the digital environment, leading to discrimination against vulnerable groups – women being less likely than men to be shown adverts for high-paid jobs, for example, or racial disparities in the predictions from algorithmic risk assessment tools in the justice system;⁸
- To society, for example through the misuse of personal data to influence political debate.⁹

This period of rapid technology development has also been accompanied by increasing concentrations of market power. Companies with access to large volumes of

information about individuals have been successful in leveraging that personal data to generate revenue. While the most prominent examples of this come from the use of personal data to enable targeted advertising for products, services, or other forms of information, access to data has provided a first-mover advantage that contributes to market concentration across the wider digital economy. While the digital economy does offer benefits to individuals, publics and policymakers are increasingly expressing concern that the benefits that come from personal data use are disjointed from the public interest.¹⁰

In many parts of the world, these technological developments have taken place alongside wider debates about the extent to which all in society are able to benefit from advances in technology and economic growth. Political shifts following the growth of populist movements in the US and Europe have prompted further concerns about the extent to which digital technologies – originally envisaged as means for democratic engagement – have created an information environment that undermines democratic discourse.¹¹ With calls to reorient the use of technologies so as to support democracy and social cohesion, governance systems are needed that align digital systems with societal values.

Addressing governments at the World Economic Forum in 1996, cyber-activist John Perry Barlow claimed in his ‘*declaration of the independence of cyberspace*’ that “cyberspace does not lie within your borders”, suggesting governments “have no moral right to rule us nor do you possess any methods of enforcement we have true reason to fear”.¹² In the early days of the growing digital economy, statements such as this fed a techno-centric narrative that argued the internet – or online activities – were beyond the reach of governments or governance. However, as digital systems become foundational to daily activities, as the vulnerabilities they create become clearer, and as publics and policymakers question who is benefitting from technological advances, this ‘wild west’ mentality seems increasingly outdated. In its place arrive new questions about power and asymmetries of power: who really controls the digital world, and how can governance help share the benefits of digital technologies across society?

Learning from this recent history, and as the disruptive potential of technologies like AI, cryptocurrencies and quantum computing become clearer, policymakers are seeking to create governance systems that allow freedom to innovate and pursue research, within an environment that pre-empts and prevents the harms that may follow. These approaches frequently seek to allow the innovation that has come from market development, in a framework that asserts democratic values.

3.2 Current legal structures and data rights

If recent history has shown the vulnerabilities created by new patterns of data use, it has also demonstrated the ability of governments, publics, civil society and industry to endeavour to mitigate these potential harms.

The last five years have seen governments across the world put in place national strategies to support the development of AI technologies and their ethical deployment. Amongst the common elements in many of these strategies are data ethics initiatives, pursued with the aim of improving the trustworthiness of AI technologies, for example:

- Germany's Data Ethics Commission was set up by the Federal Government in 2018 with a mandate "to develop ethical benchmarks and guidelines as well as specific recommendations for action, aiming at protecting the individual, preserving social cohesion, and safeguarding and promoting prosperity in the information age". The Commission recently made recommendations aimed at governing both the digital economy and AI technologies.¹³
- The UK's Centre for Data Ethics and Innovation, an advisory body established by the UK Government "to connect policymakers, industry, civil society, and the public to develop the right governance regime for data-driven technologies".¹⁴ Its first reports on bias and online targeting have sought to inform government policy development in these areas.
- Picking up the recommendations of Cedric Villani's report '*For a meaningful artificial intelligence*', France's 2018 AI strategy seeks to promote a data policy regime that encourages data sharing in the public interest, while enforcing a right to data portability.¹⁵
- At EU level, high level groups on AI and data ethics have spent recent years advising the European Commission on its approach to data and AI, with recent draft strategies on both these areas noting key areas of ethical challenge.¹⁶

Sitting alongside these strategies, there are domains in which 'hard law' sets the bounds of technology use. Legislation on net neutrality, copyright, data privacy and the use of personal data, cybersecurity, and more has begun to define what is and is not acceptable in the digital environment. These current legal frameworks provide a constellation of data rights, with different kinds of data giving rise to different kinds of rights in different jurisdictions.

Protecting vulnerable individuals from misuse of power is central to good governance. In the EU, the General Data Protection Regulation, defines a range of prohibitions on discrimination based on protected characteristics.¹⁷ It also confers rights around portability, erasability, and explainability. While many of these principles are being replicated around the world – in California's Data Freedom Act, for example – the depth of many of these provisions have not yet been tested. Further regulatory developments in the EU are expected, and seem likely to focus on the impacts of AI-enabled innovation, using assessments that take into account the risks and benefits of different applications, and recent lessons about the interactions that arise between technology advances and economic structures.¹⁸

Recent activities by regulatory bodies in many jurisdictions also signal a willingness to assertively intervene against undesirable use of data. In the US, the Federal

Trade Commission has leveraged large fines against Facebook for privacy violation¹⁹, has barred developers from selling apps that monitor consumers' mobile phone devices – so-called stalking apps – unless they “take certain steps to ensure the apps will only be used for legitimate purposes”,²⁰ and is active in reviewing anti-competitive behaviour in the sector.²¹ In the UK, the Information Commissioner's Office has similarly issued large fines against Facebook for privacy issues in handling user data, while the UK Government is currently considering how to act on recommendations made by a review of competition in digital markets.²² The European Commission is also examining data practices by large technology companies, and their implications for competition policy, with the EU's data strategy seeking to support innovation in the European technology industry and Europe's technological ‘sovereignty’.²³

These policy developments have been accompanied by an expanding pool of ethics codes and principles from the private sector and civil society. Many of these cluster around similar social and ethical issues, calling for action to increase transparency or explainability, to avoid bias or unfairness in data use and AI, to enhance privacy and security, to embed sustainability practices, and to take steps to mitigate the risks automation might pose to stable employment.²⁴

Together, these interventions seem to be questioning whether choices made by the market lead to a desirable mix of public and private interest. Such regulatory interest seems likely to be sustained over the coming years, with policymakers across the world looking for mechanisms to support data-enabled innovation, while managing the risks it creates.

As this resurgence of policy and legal interest gains pace, there is a growing movement to orient the outcomes of innovation towards beneficial societal outcomes – to ensure that technology both follows and fosters democracy. With broad consensus on the areas of concern associated with data use and AI, the challenge now is to move from these principles to actions that connect conversations about data sharing to the enforcement of individual rights.

3.3 The changing technology environment

At the same time, technologies are advancing at pace, giving rise to complex patterns of data use and decision-making. In this complex environment, data collected for one purpose can be rapidly repurposed or shared in ways that are opaque or unanticipated at the point of data collection. A 2018 study²⁵ of almost one million widely available apps found that most of those apps contained third-party tracking systems. Moreover, one in five of those apps shared data with more than twenty third parties, this data ranging from user age or gender to location details. Further analysis of these data transfer patterns showed a large number of data transfers to a handful of technology companies.²⁶ The complexity of these patterns of data exchange and aggregation mean it would be challenging for any individual to understand the destination of the data they yield to any app, creating an asymmetry in knowledge about data use. This

growing complexity of data processing compounds the limitations of consent-based models of data governance, which have been well-characterised elsewhere.²⁷

In this environment, seemingly insignificant decisions made about an individual in one area can give rise to complex effects across networks, as the outputs of different digital systems feed into – and out of – each other, and as individuals and technology interact. Any individual fact learned about an individual might be inconsequential, but – taken together, over time – the detailed picture of daily life that emerges can have a significant impact.

Existing policy frameworks are not necessarily well-placed to manage these network effects. They present different circumstances to those envisaged in the early stages of drafting the GDPR in the 1990s, where policymakers were primarily concerned with the use of data to inform decision-making in areas that might have a significant personal or social impact. The ‘first mover advantage’ that comes from having access to large volumes of data about individuals, meanwhile, favours further centralisation of data, as its aggregation enhances insights and economic benefit.

For some, these vulnerabilities foster a sense of diminished agency in the digital environment. Individuals lack power to influence the terms of data use – either because of a lack of knowledge about what choices are being made, or a lack of bargaining power in transactions – while also having their quotidian choices invisibly shaped by data-enabled systems against which there is no clear response.

Calls for returning ownership of data to individuals is one response to these challenges. However, not only is ownership unlikely to provide the level of control over the use of data that many are seeking, it is also a poor response to the vulnerabilities that are at stake.

With technology changing at pace, and complex patterns of data use and decision-making giving rise to unanticipated consequences, legislators face challenges in designing legal frameworks that allow technological progress to keep in touch with evolving socio-cultural values and expectations. What type of governance system would be best suited to a situation in which individual decisions have cumulative, unanticipated impacts?

3.4 Bridging the gaps: a democratic model for data governance?

As data governance finds itself at the heart of continuing efforts to articulate (and contest) social and political objectives, the rights granted by regulatory instruments become important tools to set limits on acceptable uses of data. However, their exercise alone is unlikely to be sufficient to give citizens a voice in shaping these data-reliant futures.

New forms of democratic governance are needed to reassert fundamental democratic values, creating a system that supports human dignity and fosters democratic representation. This requires fresh governance approaches that can bridge the gap between the aspiration to share data to achieve social and economic benefit, and concerns about protecting individual rights in data use. One approach to bridging this gap is the creation of new forms of data sharing arrangements that leverage the power that comes from aggregating data to open the way to new, bottom-up governance frameworks.

4 Commons, cooperatives, and counter-power

4.1 Mutualisation as a tool to counter power asymmetries

Data becomes valuable in aggregate. While data about an individual has limited use, collection and analysis of data about large numbers of individuals yields significant economic and social value – and power. An environment where a small number of actors have access to – or control of – this aggregated data, is one of asymmetric power, in which any single individual has limited scope to influence the terms of data use. Collective action, however, could provide a counterbalancing force.

History gives numerous examples of the ways in which combining resources can enable individuals to exert influence in systems dominated by powerful interests. In the 19th century, for example, the right to vote in the UK was conditional upon land ownership, and such ownership was available only to those with economic and social resources. Land societies were established as a means of countering this inequality. Individuals pooled their resources in a land society to collectively buy a plot of land, which was then divided between the society's membership, giving each member a right to vote. This form of mutualisation therefore gave a political voice to individuals that were otherwise disenfranchised.²⁸

Inspired by this history, it is possible to envisage governance mechanisms that seek to promote collective action. One such mechanism comes in the form of data trusts: by pooling data – or data rights – individuals would be better placed to acquire a political and economic voice in the digital economy. At stake is not the right to vote, but the ability to influence decisions about how data is used, and for what purpose.

4.2 The emergence of data trusts as a governance tool

Trusts are a legal agreement under which one party (the trustee) manages an asset or object for the benefit of another (the beneficiary).²⁹ A data trust is a mechanism to secure independent stewardship of data use under the framework of trust law.³⁰ The trust creates an intermediary layer between data subjects and controllers, with individuals that invest their data rights in a trust tasking trustees with making decisions about data use on their behalf. The ways in which data in the trust is used would depend on the terms of that trust.

Core to the functioning of a data trust are the fiduciary responsibilities trust law creates. These impose a duty of undivided loyalty that require those that lead the trust to act in the interests of its beneficiaries. These responsibilities act as a strong safeguard that sets data trusts apart from data access agreements based on contractual or corporate frameworks.

By pooling data within a trust, individuals and collectives can wield the collective power of data to exert influence over how it is used. Trusts could become powerful actors that are better placed to influence the terms and conditions of data use than any individual.

The role of the trustee sits at the heart of this mechanism, taking on significant responsibilities on behalf of the trust's members. Not only would trustees need to be mandated to exercise such rights, they would also need a set of professional skills to ensure the decisions they make are soundly-based.³¹ In the same way that previous centuries saw professionalisation of medical and legal practitioners to manage the vulnerabilities at play in those interactions, data trustees could become a new profession for the 21st century.

Data trusts would not need to be built according to a single model: some might be generalist, others built to focus on data relevant for a specific purpose; different trusts might offer different levels of participation or consultation with its members; or there could be centralised or decentralised approaches to managing the data in the trust: a trust need not hold or gather the data. By building an ecosystem of data trusts – each with different approaches to data use – individuals could select a trust that best reflects their aspirations and attitudes to risk. Individuals would be able to ‘shop around’ these different trusts, finding one that reflects their desired mix of risks and responsibilities.³²

These trusts would complement existing legal and regulatory frameworks that define the rights an individual has over how data about them is used. Instead of relying on ‘one size fits all’ regulatory approaches to setting the boundaries of data use, each trust would define its own approach to data management, taking into account the aspirations and interests of its members. In this way, trusts could offer a way of aligning an individual's values with the way their data is used.

Since gaining public prominence in 2016,³³ the ideas behind data trusts have gained traction in a variety of policy communities. The UK's Hall-Pesenti review of AI recommended that the UK Government establish trusts to promote trustworthy data sharing,³⁴ the Canadian Government's Digital Charter has recommended the creation of trusts for similar purposes,³⁵ and Germany's Data Ethics Commission has recommended that further investments be made in research and development to create data trust schemes.³⁶

With this growing interest in novel data governance frameworks, other types of data sharing institutions have emerged each with different benefits and limitations:³⁷

- Public databanks – data management institutions run by a public sector entity – that provide a publicly-accountable means of managing public data assets to deliver goods or services. While this form of institution might be able to take action to reduce the vulnerabilities associated with data use, they offer limited scope for individuals to assert how data about them is used.
- Data cooperatives that provide a means of organising data pooled from individuals or companies for a particular purpose. While offering a means for groups of individuals to promote the use of certain types of data, these lack the fiduciary safeguards inherent in reliance on trust law, since coops will be based on contractual or corporate structures. The latter structures may of course include terms that seek to prevent undesirable forms of data use, yet they will not have the same safeguards as those available under a trust structure.
- Contractual frameworks that define terms of use for data shared in specific circumstances are encountered by many people in the terms and conditions associated with data agreements. These consent-based approaches offer limited – if any – scope for individuals to influence the terms of data use and tend not to be well-suited to managing individual vulnerabilities. Horizontal data sharing agreements – another form of contractual framework that set in place access agreements between companies – can increase corporate confidence in enabling data use by providing legal certainty about acceptable use.

These different frameworks are thus more or less well-suited to different aims.³⁸ They can variably be used to promote social benefit, to protect vulnerabilities, or monetise data to different extents, or direct data use towards specific purposes that may benefit different communities. The choice of model will depend on the objectives of the data sharing activity.

Data trusts distinguish themselves from these models not only in the level of legal safeguards they provide, but also in their ability to simultaneously pursue each of these aims, as determined by their governing documents.

5 Optimising for democracy? A data governance system that benefits humanity

The power of modern AI comes from its ability to automatically extract knowledge from large amounts of data, using the insights so created to optimise systems and make predictions. The question that now pervades debates about data governance and the use of AI is: for what is the system optimised? And who decides whether this is desirable? The challenge of data governance for the 21st century is to create conduits that bring social and ethical values into technology developments, establishing mechanisms that return agency to individuals and communities.

Data governance offers a lever to reshape the underpinnings of technology development. Governance can support data use – enabling its analysis by technologies such as AI to create economic and social value – while creating an infrastructure that aligns technology development with personal, ethical and democratic values.

Some pillars of such an infrastructure already exist. Efforts to support data use through open data movements have achieved significant success over the last decade, with data availability now at the heart of many government digital services and research efforts. These open resources have already brought widespread benefits, and continue to be deployed in innovative ways. Data access agreements in recent years have been put in place to share medical data to improve diagnosis of macular degeneration, to share environmental monitoring data to tackle the illegal wildlife trade, and sharing engineering data to help address health and safety issues.³⁹ Further success for these efforts will require sustained investment to make data accessible, to make it interoperable, and to make it safe and reliable for use.

A further pillar comes from the top-down regulations that are already in place to govern data use. These will constrain actors to prevent undesirable uses of data, and create space for individual and community data rights. However, while helping to define the scope of individual rights or terms of acceptable use, these top-down endeavours cannot alone reverse the power imbalances that pervade the digital environment. The limits of these existing approaches leave a gap in the governance environment. To fill this gap, new structures are needed that provide space for individuals to collectively influence how data about them is used. These will need to be fostered by governmental action to set policy and regulatory frameworks that help such bottom-up structures grow to fulfill their potential.

Complementing existing regulatory approaches, data trusts offer a mechanism through which individuals can assert their rights, collectively gaining a voice in decisions otherwise made by a small number of people. Crucially, these trusts can bridge the gap between the widely-shared aspiration to share data to foster the realisation of various public goods on one hand and concerns about protecting individual rights on the other. In so doing, they can facilitate collective action that promotes innovative applications of data while remedying the power asymmetries that would traditionally follow such use.

Further developing the concepts and methods in the data trusts approach will require action from policymakers, industry and civil society. Understanding and overcoming the limits of existing regulatory provisions around data portability, provenance, and erasure will be necessary in order to enhance processes by which individuals can move their data between trusts.

Measures to support a wide range of individuals and communities to engage with data trusts will also be necessary in order to ensure that their benefits and protections

are accessible by all. There is generally low levels of awareness of data use and AI technologies,⁴⁰ meaning that – in the absence of any steps to promote their use – the average level of interest in registering with a trust might be low. Interventions to support citizens to understand their data rights and raise the profile of data trusts (in a way not dissimilar to pensions-related interventions) may be necessary to complement these governance structures.

Achieving the potential of AI technologies – and unlocking the value of data – requires a data environment that supports responsible data use, empowers disenfranchised groups and protects individual rights. A collection of novel data governance tools is emerging, prompting questions about the limits of existing regulatory approaches and the structures that can best embed democratic values in technology development and use. With growing interest from governments across the world in the idea of data trusts as a tool for democratising data governance, the coming years will bring a pressing need to resolve questions such as:

- Are additional legislative measures needed to enable citizens to mandate their data rights to a data trustee, or ensure the portability of their data (and/or data rights) as and when they switch from one trust to another?
- What jurisdictional issues might arise in the development of data trusts internationally, and what forms of international cooperation might be needed to address these?
- What policies or institutions should be in place to support the professionalisation of data trustees?

References:

¹ See, for example, The Academy of Medical Sciences and the Royal Society (2019) *AI in health and care: from bench to bedside*, note of discussions at a workshop on 29 March 2019, available at <https://acmedsci.ac.uk/policy/policy-projects/artificial-intelligence-and-health>

² See, for example, European Parliament (2019) *Artificial intelligence in transport*, briefing from the Members' Research Service, available at [https://www.europarl.europa.eu/RegData/etudes/BRIE/2019/635609/EPRS_BRI\(2019\)635609_EN.pdf](https://www.europarl.europa.eu/RegData/etudes/BRIE/2019/635609/EPRS_BRI(2019)635609_EN.pdf)

³ See, for example, the potential applications explored at: <https://www.climatechange.ai/>

⁴ Lawrence, N. (2019) *From data subject to data citizen*, available at <http://inverseprobability.com/talks/notes/from-data-subject-to-data-citizen.html>

⁵ See, for example, AI Now (2019) Annual Report, available at: https://ainowinstitute.org/AI_Now_2019_Report.pdf

⁶ HM Treasury (2018) *The economic value of data: a discussion paper*, available at <https://www.gov.uk/government/publications/the-economic-value-of-data-discussion-paper>

⁷ Zheleva, E. and Getoor, L. (2009) *To join or not to join: the illusion of privacy in social networks with mixed public and private user profiles*, Proceedings of the 18th international conference on World wide web April 2009, 531–540; Kosinski M, Stillwell D, Graepel T. (2013) *Private traits and attributes are predictable from digital records of human behaviours*. PNAS, 110, 5802–5805

⁸ See, for example: Datta, A., Tschantz, M., and Datta, A. (2015) *Automated experiments on ad privacy settings*. Proceedings on Privacy Enhancing Technologies, 2015 (1), 92 – 112; MIT Tech Review (2019) *AI is sending people to jail – and getting it wrong*, available at: <https://www.technologyreview.com/s/612775/algorithms-criminal-justice-ai/>; BBC (2018) *Amazon scrapped 'sexist AI' tool*, available at: <https://www.bbc.co.uk/news/technology-45809919>

⁹ See, for example: The Guardian (2018) *Cambridge Analytica Scandal 'highlights need for AI regulation'*, available at: <https://www.theguardian.com/technology/2018/apr/16/cambridge-analytica-scandal-highlights-need-for-ai-regulation>

¹⁰ CDEI (2020) *Review of online targeting*, available at <https://www.gov.uk/government/publications/cdei-review-of-online-targeting>; Ipsos MORI (2017) *Public views of machine learning: findings from public research and engagement*, available at www.royalsociety.org/machine-learning

¹¹ Pasquale, F. (2017) *The automated public sphere*, University of Maryland Legal Studies Research Paper No. 2017-31, available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3067552

¹² The full text of this statement is available at <https://www.weforum.org/agenda/2018/02/a-declaration-of-the-independence-of-cyberspace/>

¹³ Further information about the work of the Data Ethics Commission is available at https://www.bmjbv.de/DE/Themen/FokusThemen/Datenethikkommission/Datenethikkommission_EN_node.html

¹⁴ Further information about the work of the Centre for Data Ethics and Innovation is available at <https://www.gov.uk/government/organisations/centre-for-data-ethics-and-innovation/about>

¹⁵ See: <https://uk.ambafrance.org/France-s-AI-strategy>

¹⁶ See: European Commission (2020) *A European strategy for data and AI: a European approach to excellence and trust*, available at ec.europa.eu

¹⁷ Lawrence, N. (2020) *Personal Data Trusts*, available at <http://inverseprobability.com/talks/notes/personal-data-trusts.html>

¹⁸ European Commission (2020) *A European strategy for data*, available at <https://ec.europa.eu/digital-single-market/en/policies/building-european-data-economy>

¹⁹ See, for example: FTC (2019) *FTC imposes \$5 billion penalty and sweeping new privacy restrictions on Facebook*, available at <https://www.ftc.gov/news-events/press-releases/2019/07/ftc-imposes-5-billion-penalty-sweeping-new-privacy-restrictions>

²⁰ FTC (2019) *FTC brings first case against developers of stalking apps*, available at <https://www.ftc.gov/news-events/press-releases/2019/10/ftc-brings-first-case-against-developers-stalking-apps>

²¹ FTC (2019) *FTC's Bureau of Competition launches task force to monitor technology markets*, available at: <https://www.ftc.gov/news-events/press-releases/2019/02/ftcs-bureau-competition-launches-task-force-monitor-technology>

²² HM Treasury (2020) *Budget 2020*, available at <https://www.gov.uk/government/publications/budget-2020-documents/budget-2020>

²³ European Commission (2020) *A European strategy for data*, available at <https://ec.europa.eu/digital-single-market/en/policies/building-european-data-economy>

²⁴ Some of these principles are aggregated at: <http://www.linking-ai-principles.org/principles>

²⁵ Binns, R., Lyngs, U., Van Kleek, M., Zhao, J., Libert, T. and Shadbolt, N. (2018). *Third Party Tracking in the Mobile Ecosystem*. Proceedings of the 10th International ACM Web Science Conference

²⁶ Financial Times (2019) *How smartphone apps track users and share data*, available at: <https://ig.ft.com/mobile-app-data-trackers/>

²⁷ See, for example: British Academy, TechUK and Royal Society (2018) *Data ownership, rights and controls: reaching a common understanding* Available at: <https://royalsociety.org/-/media/policy/projects/data-governance/data-ownership-rights-and-controls-October-2018.pdf>

²⁸ Financial Times (2019) *Letter: Legal instruments exist to empower us, the data subjects*. From Sylvie Delacroix and Neil Lawrence, The Alan Turing Institute, available at: <https://www.ft.com/content/33926828-16c0-11ea-9ee4-11f260415385>

²⁹ Chambers, R. (2010) *Distrust: Our Fear of Trusts in the Commercial World*, 63, 631, *Current Legal Problems*

³⁰ Delacroix, S. & Lawrence, N. (2018) *Bottom-up data Trusts: disturbing the 'one size fits all' approach to data governance*, *International Data Privacy Law*, ipz014, <https://doi.org/10.1093/idpl/ipz014>

³¹ In the EU, Article 80(1) of the GDPR makes provision for such mandates to be put in place, but in only limited circumstances.

³² A legal basis for the ability to move data between trusts could be found in the EU in the GDPR's provisions on data portability and data erasure (Articles 20 and 17).

³³ Lawrence, N. (2016) *Data trusts could allay our privacy fears*, *The Guardian*, available at: <https://www.theguardian.com/media-network/2016/jun/03/data-trusts-privacy-fears-feudalism-democracy>

³⁴ UK Government Hall-Pesenti Review (2017) *Growing the artificial intelligence industry in the UK*, available at: https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/652097/Growing_the_artificial_intelligence_industry_in_the_UK.pdf

³⁵ Government of Canada (2019) *Strengthening Privacy for the Digital Age*, available at https://www.ic.gc.ca/eic/site/062.nsf/eng/h_00107.html

³⁶ Daten Ethik Kommission (2019) *Opinion of the Data Ethics Commission*, available at <https://www.bmjv.de/>

³⁷ These are explored further in Delacroix, S., Lawrence, N. and Montgomery, J. (2020) *Selecting a data sharing structure: a value-based choice*, available at <https://datatrusts.uk/blogs/selectingdatastructures>

³⁸ Delacroix, S., Lawrence, N. and Montgomery, J. (2020) Selecting a data sharing structure: a values-based choice. Available at: <https://datatrusts.uk/blogs/selectingdatastructures>

³⁹ See, for example: Open Data Institute (2019): *Data trusts: lessons from three pilots (report)*, available at <https://theodi.org/article/odi-data-trusts-report/>

⁴⁰ Royal Society and Ipsos MORI (2017) *Machine learning – what do the public think?* Available at www.royalsociety.org/machine-learning