The Role of Technology in Governance: the Example of Privacy Enhancing Technologies

Dr Natasha McCarthy and Dr Franck Fourniol

Dr Natasha McCarthy, Head of Policy, The Royal Society
natasha.mccarthy@royalsociety.org
Dr Franck Fourniol, Policy Adviser, The Royal Society
franck.fourniol@royalsociety.org

This paper is based on the Royal Society report: Protecting privacy in practice – the current use, development and limits of Privacy Enhancing Technologies in data analysis[1]

Abstract

The collection of data, its analysis and the publication of insights from data promise a range of benefits, but can carry risks for individuals and organisations. This paper sets out a way to identify the right role for technologies in governance of data use. The paper examines the potential of Privacy Enhancing Technologies to support governance of data use, and considers their role based on their current state of development and the trajectory of technological development. This involves consideration both of how these technologies can potentially enable governments and others to unlock the value of data; and also recognition of both contingent and in principle limitations on the role of PETs in ensuring well-governed use of data.

Keywords – Privacy Enhancing Technologies; governance; data management and use; homomorphic encryption; differential privacy; trusted execution environments; secure multi-party computation; personal data stores

Background and introduction

The amount of data generated from the world around us has reached levels that were previously unimaginable. The use of data-enabled technologies promises significant benefits, from improving healthcare and treatment discovery, to better managing critical infrastructure such as transport and energy. However, the collection of data, its analysis and the publication of insights from data can all carry risks for individuals and organisations.

The British Academy and Royal Society report *Data management and use: Governance in the 21st century*,[2] published in June 2017, highlighted a series of such tensions between benefits and risk in the use of data. These tensions include:

• Using data relating to individuals and communities to provide more effective public and commercial services, while not limiting the information and choices available.

• Promoting and distributing the benefits of data use fairly across society while ensuring acceptable levels of risk for individuals and communities.

• Promoting and encouraging innovation, while ensuring that it addresses societal needs and reflects public interest.

• Making use of the data gathered through daily interaction to provide more efficient services and security, while respecting the presence of spheres of privacy.

The Royal Society has continued this work on the governance of data use by considering the role of Privacy

[1] The Royal Society (2019) Protecting privacy in practice – the current use, development and limits of Privacy Enhancing Technologies in data analysis (see https://royalsociety.org/topics-policy/projects/privacy-enhancing-technologies/, accessed 3 May 2019)
[2] The British Academy and the Royal Society (2017) Data management and use: Governance in the 21st Century (see https://royalsociety.org/topics-policy/projects/data-governance/, accessed 3 May 2019)

Enhancing Technologies (PETs) in enabling well-governed use of data. The Privacy Enhancing Technologies project has been looking into the role of technologies in enabling data analysis and extracting value whilst preserving personal or sensitive information.

The aims of the project were to explore the interplay between the following questions in relation to PETs:

- What are the ethical and social issues at stake?
- What is mathematically possible and what is technically feasible?
- What business models and incentive systems can deliver these technologies?

A high-level aim of the project was to assess the extent to which there is a role for technology in addressing the tensions set out in the *Data Management and Use: Governance in the 21st century* report, and to understand the specific roles and limitations of PETs in enabling data users – including those in the public sector – to steer a course through these tensions. In doing so, it explored the underlying principles, current state of development, and use cases for the following technologies, considering them within the wider context of business systems and governance frameworks influencing data use:

1) Homomorphic Encryption
2) Differential Privacy
3) Secure Multi-Party Computation
4) Trusted Execution Environments
5) Personal Data Stores

Method

Through a series of workshops with participants from academia, industry, government and the third sector, the Royal Society explored:

- The social and business trends influencing the ways that data is used and managed
- Potential business models supporting the uptake of PETs
- The stage of development of a range of PETs and their areas of application

The outputs of these workshops, alongside a series of use cases, have been brought together in a policy report setting out the roles, uses and limitations of PETs, published in March 2019: *Protecting privacy in practice – the current use, development and limits of Privacy Enhancing Technologies in data analysis* (Ibid).

The role of PETs in data governance – how can PETs enable well-governed use of data?

Navigating the tensions, set out above and in the *Data Management and Use: Governance in the 21st century* report, requires appropriate governance mechanisms, from codes of conduct and ethics to regulation. However, in some cases, technological solutions can help diffuse dilemmas between making use of data and protecting both the individuals and organisations that generate or are subjects within datasets. PETs as a category comprises a broad suite of technologies and approaches – from a piece of tape masking a webcam to advanced cryptographic techniques. While some are focused on protecting private communications, our report explored a subset of five PETs identified during the scoping of the project as being particularly promising to enable privacy-aware data collection, analysis and dissemination of results.

The key question of this paper is whether, according to the current state of development and the trajectory of technological development, we can utilise PETs in addressing social and ethical tensions in data use, and thereby use them as tools of governance. This will involve consideration both of how these technologies can potentially enable governments and others to unlock the value of data; while also recognizing both contingent and in principle limitations on the role of PETs in ensuring well-governed use of data.

The technologies

There is currently no technology that is applicable to every single situation of privacy-preserving data analysis. Different PETs can be used to achieve distinct aims, such as:

- securely providing access to private datasets
- enabling joint analysis on private data held by several organisations
- securely out-sourcing to the cloud computations on private data
- de-centralising services that rely on user data

One way to understand the technologies that we set out in the report and their potential role in governance is by relating them to the social and ethical tensions set out in

Data Management and Use, which they might help to resolve. To note again that we were primarily looking at privacy protection for big data analysis – but with the exception of personal data stores which present a slightly different approach.

Tension 1: **Making use of the data** gathered through daily interaction to provide more efficient services and security, whilst respecting the presence of **spheres of privacy.**
Tension 2: Providing ways to **exercise reasonable control over data** relating to individuals whilst **encouraging data sharing** for private and public benefit.

Homomorphic encryption Forms of this cryptographic approach essentially enable analysis of data while protecting sensitive information, by making it possible to compute on encrypted data without deciphering it.

Homomorphic encryption can be used to analyse data in circumstances where all or part of the computational environment is not trusted, and sensitive data should not be accessible. It is currently applicable where the computation required is known and relatively simple. Homomorphic encryption provides confidentiality and can be used to address the risk of revealing sensitive attributes related to individuals or organisations, in a dataset or output.

Trusted Execution Environments A Trusted Execution Environment (TEE) is a secure area inside a main processor. TEEs are isolated from the rest of the system, so that the operating system cannot read the code in the TEE. However, TEEs can access memory outside. TEEs can also protect data 'at rest', when it is not being analysed, through encryption.

Like homomorphic encryption, TEEs might be used to securely outsource computations on sensitive data to the cloud. Instead of a cryptographic solution, TEEs offer a hardware-based way to ensure data and code cannot be learnt by a server to which computation is outsourced. Unlike homomorphic encryption, current TEEs are widespread and permit the computation of virtually any operations.

Tension 3: **Incentivising innovative** uses of data whilst ensuring that such data can be **traded and transferred** in mutually beneficial ways.

Multi Party Computation Secure multi-party computation (MPC) is a subfield of cryptography concerned with enabling private distributed computations. MPC protocols allow computation or analysis on combined data without the different parties revealing their own private input. In particular, it may be used when two or more parties want to carry out analyses on their combined data but, for legal or other reasons, they cannot share data with one another.

Forms of this enable different organisations to analyse and derive insights from data without pooling it together – there is commercial value here. How can companies learn from each other without giving away trade secrets?

Tension 1: **Making use of the data** gathered through daily interaction to provide more efficient services and security, whilst respecting the presence of **spheres of privacy.**
Tension 4: **Promoting and distributing the benefits** of data use fairly across society while ensuring **acceptable levels of risk** for individuals and communities.

Differential Privacy Differential privacy is slightly different in that it is not a technology per se but is an approach which enables us to set a limit to the amount of privacy risk we are willing to take in order to access data. It is achieved by adding noise to a dataset of results of a computation, and the result is a limit on the chances of an individual or organisation being identified in that data.

Tension 2: Providing ways to **exercise reasonable control over data** relating to individuals whilst encouraging data sharing for private and public benefit.

Personal Data Stores Personal Data Stores (PDS) are again not a technology as such, but they are ways of enabling better control of data so that people can access data-driven services without giving away their data to an external server.

Unlike the other four PETs covered in the report, which are tools for privacy-preserving computation, PDS are consumer-facing apps and services which can be supported by different kinds of PETs. They provide an example of one of the goals for PETs – enabling people to have more control over data.

PDSs enable a distributed system, where the data is stored and processed at the 'edge' of the system, rather than centralised. It is possible, for instance, to send machine learning algorithms to the data, rather than the data to the

algorithms. Distributing out the data and computing solves a number of issues such as the 'honeypot' issue – whereby an organisation holding millions of records constitutes a 'honeypot' that is economically attractive to hack.

These different technologies and approaches are used to varying degrees depending on their level of maturity. There are some examples of current pilots and well-established use.

Use cases: what can work in practice

How far are these technologies able to perform real-world governance functions? The field of PETs development is moving quickly, and the Royal Society report captures a moment in time where the technologies are maturing and opportunities to use these technologies are beginning to emerge. It may be that some of the technologies surveyed in our report do not achieve their promise in the near term, or that the costs of adoption prove prohibitive, or that other technologies not explored in depth might leapfrog them. However, there are a number of areas where PETs are already in use.

Beginning with **secure multi-party computation**, the first real-world application of Sharemind – which uses MPC – was the analysis of Key Performance Indicators (KPIs) for the Estonian Association of Information Technology and Telecommunications (ITL). The ITL proposed collecting certain financial metrics and analysing them to gain insights into the state of the sector. The member companies expressed concerns over the confidentiality of the metrics, as they would be handing them out to competitors.

This prompted the use of MPC, with Sharemind developing a solution that was deployed in 2011.[3] 17 participating companies acted as the input parties who uploaded their financial metrics to three computing parties with the capability to host the Sharemind platform. ITL management acted as the result party, leading the processing and dissemination of results.

Differential privacy has been put into practice by a number of organisations handling large amounts of data, to assess and limit the risk to individuals' privacy. For example, in 2017, the US Census Bureau announced that it would be using differential privacy as the privacy protection mechanism for the 2020 decennial census.[4] This is having already implemented differential privacy for other services: e.g. onTheMap, 2008, an online application developed in partnership with 50 US states, for creating workforce related maps, demographic profiles, and reports. By incorporating formal privacy protection techniques, the Census Bureau will be able to publish a specific, higher number of tables of statistics with more granular information than previously. By fixing a privacy budget for that given set of released publications, the institution can reason mathematically about the risk of disclosure of information relating to a specific individual.

In order to share NHS data securely with multiple teams, whilst maintaining as much as possible the potential usefulness of the data, NHS Digital have been using a de-identification service employing **homomorphic encryption**. For security reasons, data is de-identified in different 'pseudonymisation domains' for each different part of an organisation. Within one domain, all data with the same base value is replaced with the same 'token' (a non-identifying value). Across domains, the same base value receives different token. Usually, transferring data between domains requires to remove the encryption for the first domain and replace it with the second domain encryption. However, using consistent 'tokenisation' and partially homomorphic encryption by Privitar Publisher, it is possible to transform data items between any two domains without revealing the base value, even if they have been de-identified by two instances of the de-identification service using different encryption keys.

This methodology allows the de-identification tool set to be deployed to multiple locations across the NHS and makes any data de-identified by any tool from the de-identification tool set potentially linkable with any other data de-identified by any other tool from the tool set.

[3] Archer DW et al. (2018) From Keys to Databases – real-World Applications of Secure Multi-Party Computation (see https://eprint.iacr.org/2018/450, accessed 3 May 2019)

[4] Garfinkel SL et al. (2018) Issues Encountered Deploying Differential Privacy (see https://arxiv.org/pdf/1809.02201.pdf, accessed 3 May 2019)

Constraints: what are the technical limitations of PETs in practice?

When using PETs, there are trade-offs. Privacy engineers say that PETs incur a cost in terms of 'utility'. In the context of different technologies, the cost in utility might be of a different nature. It is also important to bear in mind that a number of these technologies are still in a research phase, and technical limitations might therefore evolve in time. For example, with differential privacy adding noise to a dataset entails a loss of some useful information so there is a cost in terms of accuracy. The first organisations that successfully implemented differential privacy have been able to do so because they are handling large amounts of data, and the effect of noise then is less severe in terms of loss of information.

In the case of PETs where computation happens on encrypted data, such as homomorphic encryption and secure multi-party computation, the main cost to utility is in terms of computation resources (time, computing power). Encryption can entail a substantial increase in data size, which can cause a major bandwidth problem.

In order to negotiate these trade-offs, users need to have a clear idea of what information or value they are trying to protect, and they need to determine the potential benefits and costs of different PETs so that systems can be optimised for this. It is, for example, important to consider the financial cost associated with enforcing a given trust model, in particular if a trusted authority needs to be appointed.

Limitations: how far can PETs deliver ethical use of data?

The key consideration in the use of PETs as a tool of governance, is that their use does not in itself automatically make an analysis legal, ethical or trustworthy. There are a range of risks posed by data analysis, some can be addressed by PETs, others not.

When making use of PETs, we might consider in particular the following kinds of risks posed by data analysis:

- How much does the analysis reveal about the whole population or group from which the data used for the analysis originated? (Note that there is a definitional issue here: some might argue that this question relates rather to fairness and discrimination, and others might point out that it also affects the privacy of individuals in the population)

- Does the analysis reveal whether someone or a specific entity is included in the dataset that was used to conduct the analysis?

- How much does an analysis reveal about sensitive attributes about specific individuals or entities in the dataset?

- To whom is information revealed and what might they do with it?

- How sensitive are the input, intermediate values and output of an analysis?

The explanations of the technologies above indicate how the technologies might address these risks. However, many ethical questions arise through the data analysis pipeline, which are not grouped within these particular kinds of questions. In assessing whether data analysis is ethical we might consider broader concerns, e.g. whether the purpose of data use is socially beneficial, whether it might result in disadvantages for certain individuals or groups, whether the data has been collected appropriately, and so on. Implementing PETs can ensure that the methods by which data is used include protection against specific privacy risks, but it has a less direct relationship to these broader ethical

concerns, and often involve certain ethical or governance questions being settled in advance of their being used.

For example, a key question in data governance is, should we be collecting this data at all in the first place? One of the things PETs may help with is to provide extra scrutiny of this question. Given that the use of a PET incurs certain costs, as set out above, there has to be a significant benefit to collecting and using data, if you intend to use PETs to ensure that analysis of that data is privacy-preserving. So one of the key questions might concern, not how we use the technologies to protect the data through collection and analysis, but rather focus on whether there is any pragmatic (and also ethical) purpose for getting the data at all.

Moving on to data analysis, use of some of the PETs actively require certain ethical and social questions to have been settled in advance of their being used. For example, when using Differential Privacy a 'privacy budget' has to be set which establishes the acceptable risk that an individual might be identifiable in the output of an analysis – e.g. if a specific person could be identified through the statistics published by the US Census Bureau in the example above. This cannot be set by technology alone and is itself an act of governance. Legal requirements might mean that organisations will have to observe a 'minimum' privacy aim, and guidance from regulators might help improve understanding of this.

Furthermore, PETs have the potential to actively enable unethical use of data, by virtue of potentially enabling computation in private. For example, PETs such as MPC might enable companies to misbehave and to form collusions, e.g. to set prices or other aspects; some research is addressing this threat.[5]

All of this means that there are, in principle, limits to the governance role of these technologies, or requirements that certain governance questions be addressed separately to the technologies being utilised. They are not a 'solution' to a 'problem' posed by the need to balance risks and benefits of data-enabled technologies but they are a tool to be used to put governance in place. As such need skilled users and the right business and governance environment to achieve the desired outcome.

The road to adoption: steps in enabling appropriate uptake and development of PETs

How do we create these skilled users, and good environments? Our report made a number of recommendations about the route to appropriate adoption of PETs.

Accelerate research and encourage development and adoption

First, there is a challenge of developing the technologies to address limits created by the current stage of technology development. One way of focusing research is by funders, government, industry and the third sector working together to articulate and support the development of cross-sector research challenges. Alongside providing continued support for fundamental research on PETs this can potentially support development of PETs so that they address particular governance needs.

Government can be an important early adopter, using PETs and being open about their use so that others can learn from their experience. Government departments should consider what existing processing might be performed more safely with PETs and how PETs could unlock new opportunities for data analysis, including opening up the analysis of

[5] Alwen J. et al (2009) Collusion-Free Multiparty Computation in the Mediated Model. In: Halevi S. (eds) Advances in Cryptology - CRYPTO 2009. CRYPTO 2009. Lecture Notes in Computer Science, vol 5677. Springer, Berlin, Heidelberg

sensitive datasets to a wider pool of experts whilst fully addressing privacy and confidentiality concerns.

Supporting intelligent users

To enable wider and well-informed adoption, we need to create and support a community of intelligent users – who know what the technologies can and cannot deliver both technically and in terms of meeting ethical requirements.

There is a need for government, public bodies and regulators to raise awareness further and provide guidelines about how PETs can mitigate privacy risks and address regulations such as GDPR. For example, the Information Commissioner's office (ICO) should provide guidance about the use of suitably mature PETs to help UK organisations minimise risks to data protection, and this should be part of the ICO's Data Protection Impact Assessment guidelines. Such guidance would need to cover how PETs fit within an organisation's overall data governance infrastructure, since the use of PETs in isolation is unlikely to be sufficient.

To give public sector organisations in particular the level of expertise and assurance they need to implement new technological applications, a centralised approach to due diligence would be beneficial and help assure quality across the board.

The National Cyber Security Centre should act as a source of advice and guidance on the use of suitably mature PETs, as part of a network of expert organisations. Such a network of expertise would support the development and evolution of best practices and also provide access to advice on specific cases of data use or sharing. Ultimately, this could also serve as a point of engagement for academics and industry bodies working in the space and provide a portal from which private sector organisations interested in learning about PETs could access information on existing case studies.

Standards and kitemarks are needed for quality assurance and to increase 'buyer confidence' in PETs. Currently privacy standards are unclear and guidelines are scarce. Even though there is a lot of research on standards and processes, currently they are not mature enough for cross-sector agreement on best practice.

An integrated approach to governance

Regulators and civil society need to consider how PETs could become part of the data stewardship infrastructure, underpinning governance tools such as 'data trusts' and other initiatives for the governance of data use. In the UK this means the Department for Digital, Culture, Media and Sport (DCMS), the Centre for Data Ethics and Innovation (CDEI), office for AI, and other bodies coming together to discuss the right role for these technologies as they develop further. If we are to unlock the benefits of data for policy, balancing the use of technologies with the right organizational structures, institutions, and skilled users is important to ensure human flourishing in a data-enabled society.

Acknowledgements

The authors are grateful to the Chair and Working Group of experts who have led the Royal Society's Privacy Enhancing Technologies project and overseen the production of the report on which this article is based: Professor Alison Noble FRS FREng (Chair), Guy Cohen, Professor Jon Crowcroft FRS FREng, Dr Adria Gascon, Marion Oswald and Professor Angela Sasse FREng.