The Anonymisation Decision-making Framework

Elaine Mackey, Mark Elliot, Kieron O'Hara

1. Introduction

The need for well-thought-out anonymisation has never been more acute. In the current data environment, increasing numbers of large and linked datasets are created about people. At the same time there are political, social and economic drivers encouraging greater sharing of anonymised data. Add into the mix that anonymisation is a complex process and then it is perhaps not surprising that the drive to share data has led to some ill thought out and poorly-anonymised data publications including the Netflix[1], AOL[2] and New York taxi[3] cases, underlining how important it is to carry out anonymisation properly and what can happen if you do not.

In this paper, we introduce a framework for undertaking well thought out anonymisation: the Anonymisation Decision-making Framework (ADF). The ADF attempts to unify the technical, legal, social and ethical aspects of anonymisation to provide a comprehensive guide to doing anonymisation in practice[4].

In the first section of this paper we outline the thinking behind the ADF. We then go on to consider anonymisation within the framework of UK data protection law and explain how the ADF can help practitioners meet their obligations under it. In the second section we provide a brief overview of the ADF.

2. Principles underpinning the ADF

There are five principles on which the ADF is founded. They inform how we: (i) describe the problem space and (ii) establish ways to address it. Those five principles are:

1. *You cannot decide whether data are safe to share or not by examining the data alone.*

    This principle is fundamental to our understanding of re-identification risk and for finding ways of addressing it. What is alluded to here is that re-identification risk does not reside in the properties of a dataset alone but rather arises from the interaction between that dataset, people, other data and the structures that shape those interaction such as security systems,

[1] See CNN Money (2010) http://tinyurl.com/CNN-BREACHES
[2] See Arrington (2006) http://tinyurl.com/AOL-SEARCH-BREACH
[3] See Atokar (2014) http://tinyurl.com/NYC-TAXI-BREACH
[4] The ADF is presented in full in UKAN's book *The Anonymisation Decision-making Framework* which can be downloaded free from http://ukanon.net/ukan-resources/ukan-decision-making-framework/

governance practices, legislative frameworks, national policies on data sharing etc. The core features – people, other data and structure make up what we term as the data environment. It is our view that anonymisation is a heavily context-dependent process which requires consideration of data and its environment as a total system. This approach we called a *data situation approach*. It is underpinned by another core concept: *functional anonymization* which asserts that data can only be determined as anonymised or not in relation to their environment. Thus the critical question for practitioners is not 'is this data anonymised' but 'is this data anonymised in environment a, in environment b' and so on.

2. *But you still need to examine the data.*
   You need to know your data and establish its risk profile. Once you have developed a risk profile for your data you are in a good position to start thinking through in a formal way how the core features in the data environment (people, other data, and structure) may interact with your data to increase or negate the risk of re-identification.

3. *Anonymisation is a process to produce safe data but it only makes sense if what you are producing is safe useful data.*
   This principle is one side of the coin to that of principle 4. Anonymisation is a process not a simple end state, it should be understood as inseparable from its purpose to share data. There is little point, after all, in releasing data that do not represent whatever it is they are meant to represent.

4. *Zero risk is not a realistic possibility if you are to produce useful data.*
   This is fundamental; anonymisation does not guarantee zero risk of re-identification. Indeed the DPA (1998) and (it would seem the) GDPR recognises this, there is no requirement to remove risk entirely but rather what is required is that those sharing data mitigate the risk of re-identification until it is **remote** (please see UK: Information Commissioner's Office 2012). Anonymisation then is best understood as a risk management process. Accepting that there is a residual risk in all useful data inevitably puts you in the realms of balancing risk and utility. But this is the stuff of modern life – the trade-off of individual and societal level benefits against individual and societal level risks.

5. *The measures you put in place to manage re-identification risk should be proportional to the risk and its likely impact.*
   Balancing risk and utility requires proportionate decisions making about your data and who should have access to it and how. The ADF can help you in establishing sound and proportional measures to produce safe useful data.

2.1 How do we describe the problem space

The ADF is underpinned by a relatively new way of thinking about the problem of re-identification risk. It posits that you must look at both the data to be shared and the data environment to ascertain realistic measures of risk, which as we have said is called the data situation approach.

Perhaps it seems obvious that the environment in which data are to be shared is important, but for many years the data confidentiality field has focused almost

exclusively on the data themselves. Thus re-identification risk was seen as originating from and largely contained within the data to be shared and as a consequence, researchers and practitioners rarely looked beyond the statistical properties of that data. With a few notable exceptions (e.g. Duncan & Lambert 1989, Elliot and Dale 1999 and Reiter 2005) they have not concerned themselves with issues such as how or why a re-identification might happen, or what skills, knowledge, or other data a person would require to ensure his or her intrusion attempt was a success. This has meant that the models they built to assess risk, whilst statistically sophisticated, have at best been based on assumptions about the data context and at worst totally detached from any real-world considerations. The data situation approach attempts to remedy this, by broadening our understanding of the problem space to include the actions of key agents, other data within the environment and previously-neglected considerations such as the importance of governance processes and ethical decision-making. The basic premise is that you cannot guard against the threat of re-identification unless you have a clear idea of what it is you are guarding against and this requires considering both data and environment.

1.2 Anonymisation, the law and the ADF

Anonymisation, as we have described thus far as a top level concept, is a complex decision-making process of risk management. At a narrower level, it is a process that enables data custodians to transform personal data in such a way as to render it anonymised for the purpose of sharing it.

In the UK, the law most relevant to personal data and their anonymisation is the 1998 Data Protection Act (the DPA)[5]. In 2018, the new European Data Protection Regulation will come into force[6].

The regulation will be a significant change not just in how it is enforced (as a directly binding legislative act that must be applied in its entirety across the EU) but also in content. For example, in the new regulation:

- The definition of what is personal data has been expanded to include location data, genetic data and online identifiers.
- There are changes to the rules for obtaining valid consent.
- Pseudonymisation is introduced as a privacy enhancing technique.
- Data protection/privacy by design will be an obligation.
- Notification of breaches will be mandatory.
- The data processors' responsibilities are increased.

[5] Other legislation that pertains to particular datasets and data sources such as the *Statistical Registration and Services Act* (2007) for official statistics, the *Commissioners for Revenue and Customs Act* (2005) for HMRC data and the *Census (Confidentiality) Act* (1991) for UK Census data, is also pertinent for anonymisation in the UK, and approaches it via the notion of *identifiability* which is central to the DPA's concept of personal data.

[6] If the new regulation comes into force whilst the UK is still an EU member state, it will apply to UK law. If the UK leaves the EU, it is the view of the UK Information Commissioner that the UK will have to have legislation that is comparable with the rest of the EU for trading purposes, (Killen, 2016)

- Data subjects are given enhanced rights including the right to be forgotten, the right to object to automated decision-making and data portability.

The textual definition of what are personal data and the functional role of anonymisation in data protection however appear to be unaltered by the new regulation.

Personal data is taken to mean, data that: relates to a living person who can be identified either:

(a) directly from those data, or
(b) indirectly from those data and other information.

It is the second of these conditions and in particular the terms 'indirectly and other information' that gives anonymisation its inherent complexity. This is because:

(i) It is not always apparent (once formal identifiers such as name and address, national insurance number etc. are removed or obscured) what variables or combinations of variables may be identifying.

(ii) Even with the requirement of the *means reasonably likely* test in the DPA (and new regulation) there has been, to date, little in the way of concrete guidance for practitioners on how to think about and analyse their data in relation to its environment(s).

It is perhaps not surprising then that we see: data shares that are shown to be disclosive, see for example Ohm 2010; Narayanan and Shmatikov 2010, or cases contested through data protection regulators and the courts about whether or not data can be shared (in an anonymised form) e.g. PACE[7]; or data providers behaving risk adverse and not sharing data when there is a legal basis for doing so. Certainly data environments are multifaceted and complex and, in the case of open data environments in particular, not easily defined. Added to this there has been little work beyond that of the University of Manchester's Data Environment Analyse (initially through the DEAS[8] and later UKAN[9]) on what additional information might be available through public sources such as Registers and the internet and restricted information sources such as organisational databases. What has been lacking and the ADF provides is a formal mechanism for assessing the disclosure risk of data products in the context of the environment(s) in which they are held and shared.

Before we go on to introduce the ADF there is one further point we would like to make in relation to how the framework can help practitioners process their data in accordance with their obligations under the new regulation. The regulation expands the obligation on data controllers to implement appropriate technical and organisational measures to protect personal data to ensure that privacy and the protection of data are no longer an after-thought (see for example in Recitals 78 and 108, and Article 47). It does this through the concept of 'data protection by design', and although its meaning is not entirely clear the notion of privacy by design (often

[7] For further information http://blogs.bmj.com/bmj/2016/09/22/peter-white-et-al-releasing-patient-data-from-the-pace-trial-for-chronic-fatigue-syndrome/

[8] The Data Environment Analysis Service (DEAS) was a collaboration between Manchester University and the Office for National Statistics. As part of its work it identified, collated and categorised external sources of information for informing disclosure scenarios. See for example, Elliot et al 2010.

[9] UKAN is a collaborative network set up to provide best practice guidance on anonymization to anyone who holds personal data and needs to share it.

referred to by commentators in relation to data protection by design, see for example ICO https://ico.org.uk/for-organisations/data-protection-reform/overview-of-the-gdpr/accountability-and-governance/) is well documented. It is a concept introduced by Cavoukian in the 90's and is characterised by seven principles, namely that one's approach should be:

1. Proactive, not reactive, preventative, not remedial.
2. Privacy as the default setting.
3. Privacy embedded into design.
4. Full functionality. Positive sum, not zero sum.
5. End-to-end security. Full lifecycle protection.
6. Visibility and transparency.
7. Respect for user privacy. Keep it user-centric.

The Anonymisation Decision-making Framework can facilitate a privacy by design approach to anonymisation through its data situation approach, which encourages the practitioner to: plan their (intended) anonymisation activities; plan for the unexpected and potential consequences of sharing data; create a clear audit trail of their anonymisation activities; be transparency; respect data subjects; achieve optimal data safety and utility.

2 The Anonymisation Decision-making Framework

The ADF is not an algorithm; it is an approach whose value depends on the extent of the knowledge and skills that the practitioner brings to it. Nor is it a simple check list that can be run through in a linear fashion. All the important considerations are there but for the practitioner using the framework they will need to think how each relate to and impact on the others in a given data situation.
The framework is made up of the ten components. These encompass three core anonymisation activities:

- *A data situation audit:* a framing tool for identifying those issues relevant to a proposed data share or release (covered by components 1-5).
- *Risk analysis and control:* the technical processes needed to assess and manage the disclosure risk associated with a data situation (covered by components 6-7).
- *Impact management:* measures to manage the (expected or potential) consequences of a share (covered by components 8-10).

For a full exposition of the components see Elliot et al (2016) here we provide a brief summary.

Component 1: Describe your (intended) data situation

In this component we set out what we mean by a data situation, which essentially describes the relationship between some data and their environment. It is an important conceptual tool for helping practitioners to explain and visualise their own particular data situation. The parameters of a data situation can be determined by mapping the intentional flow of data from one environment to another.

Data situations can be *static* or *dynamic*. A static data situation is where there is no movement of data between environments; a dynamic data situation is where there is such movement. By definition all data shares take place within dynamic data situations in which data are intentionally moved from one environment to another. A dynamic data situation might be relatively straightforward involving the movement of data from just one environment to another environment. Often though, it is more complex involving multiple environments.

Let us consider this further using an example. Imagine that PubT (a franchised public transport provider) collects personal data from its customers relating to public transport usage. PubT plans to share an anonymised version of the data with the Local Authority of Barsetshire which wants to use it to help with infrastructure planning. PubT anonymises the data by removing the direct identifiers, i.e. the customers' names and addresses, and by aggregating the detail on several key variables. However, it leaves some key variables – which are of particular interest to Barsetshire – unchanged. We shall call this environment 1.

Barsetshire signs a contract with PubT which (i) enables it to analyse the data, (ii) proscribes it from sharing any part of the data without prior agreement and from holding the data for longer than a specified time period and (iii) ensures it is kept securely. After this contract is signed, the anonymised dataset is passed to Barsetshire, this is environment 2.

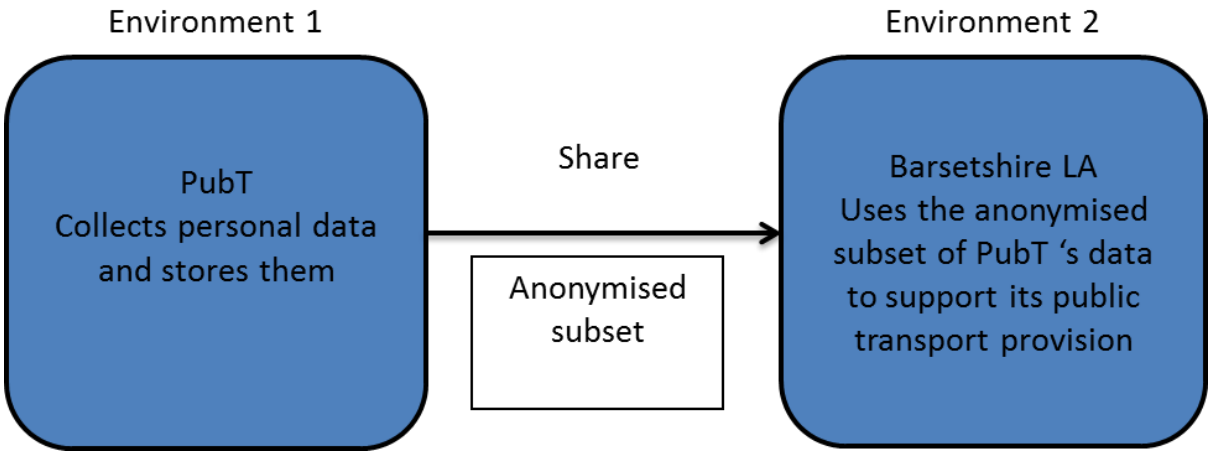

Figure 1: data flow between environments

Figure 1 illustrates the intentional movement of data from environment 1 to environment 2. The data flow between PubT and Barsetshire defines the parameters of the data situation. By using a contract to stipulate how the data can be processed and accessed, PubT is placing controls on governance processes and infrastructure within environment 2 and thereby controls the disclosure risk associated with the data situation. The (anonymised) data within Barsetshire's environment is considered low risk even though it contains some detailed key variables. This is because the environment is restricted – few people have access to the data and their use of the data is clearly defined and managed. Conversely, in this scenario the data

is unlikely to be considered safe within a less restricted environment, such as an open access environment, because no such controls would be in operation. This may seem obvious, but failure to understand the basic point that data releases need to be appropriate to their release environment is the primary cause of the well-publicised examples of poorly anonymised datasets.

Component 2: Understand your legal responsibilities

The movement of data across multiple environments can complicate the question of who is responsible for it. In this component we address the question of whether you are a data controller, processor or user by considering:

- The status of the data in each environment (is it personal data).
- The provenance of the data (who collected the data, who determined the legal basis for collection etc.).
- The enabling conditions for the share (how is the processing fair and lawful).
- The mechanism for the share (e.g. a data sharing agreement, contract, licensing agreement).

Component 3: Know your data

Here, we describe the factors you need to take account of to carry out a risk profile: a top level assessment of your data. We first explore how to determine the status of data and then consider how disclosure risk may be affected by the following:

- Data type: The form of your data, e.g. statistics or text; the level of information e.g. microdata or aggregated.
- Variable types: The nature of each variable within the data: direct identifiers, indirect identifiers or targets.
- Dataset properties: The top level properties of the dataset, e.g. its age, quality, etc.

Component 4: Understand the use case

Establishing the use case for the data will help you think about the mapping of what data is required/beneficial to the end user and what you are able provide. In this component we explain how you can go about doing this, by:

- Clarifying your reason for sharing or releasing your data
- Identifying the user groups who may wish to access your data
- Establishing how those accessing your data might want to use it

Component 5: Meet your ethical obligations

Here we considered how you can meet your ethical obligations whilst maximising the value of your anonymised data. We consider the issues of:

- Consent and awareness and how these relate to the ethical load of your data situation.
- Ameliorating actions such as transparency and stakeholder engagement.
- The importance of good governance.

Component 6: Identify the processes you will need to assess disclosure risk

In this component we introduce a four-part process for assessing disclosure risk. The first two procedures are always necessary, while the third and fourth may or may not be required depending on the conclusions drawn after conducting the first two.

1. Incorporation of your top level assessment (risk profile undertaken in component 2) to produce an initial specification.

2. An analysis to establish relevant plausible scenarios for your data situation. When you undertake a scenario analysis, you are essentially considering the how, who and why of a potential breach.

3. Data analytical approaches. You will use data analytical methods to estimate risk given the scenarios that you have developed under procedure 2.

4. Penetration testing, this involves validating assumptions made in procedure 2 by simulating attacks using 'friendly' intruders. The ICO recommends carrying out a motivated intruder test as part of a practical assessment of a dataset's risk. This can be both informative and good practice but takes skill and expertise as well as time and resources.

Component 7: Identify the disclosure control processes that are relevant to your data situation

Disclosure control processes essentially attend to either or both of the two elements of your data situation: the data and their environment. In this component we explain how if your risk analysis in component 6 suggests that you need stronger controls then you have two (non-exclusive) choices to:

- Change the data (specification)
- Reconfigure the data environment

Component 8: Identify who your stakeholders are and plan how you will communicate

In this component we outline how effective communication can help build trust and credibility, both of which are critical to difficult situations where you may need to be heard, understood and believed. The key point being you will be better placed to manage the impact of a disclosure if you and your stakeholders have developed a good working relationship.

Component 9: Plan what happens next once you have shared or released the data

Having carried out an anonymised data share, we consider what you subsequently need to do in respect of those data; specifically we detail what you can do to monitor the environment into which you release, by for example:

- Keeping a register of all shares and data disseminations.
- Comparing planned shares against previous shares to prevent the risk of data linkage.

Component 10: Plan what you will do if things go wrong

Sometimes, even when you follow all the recommended advice, things can go wrong. Here we describe what you can do to help deal with a disclosure in the rare event that one occurs.

3 Concluding remarks

The anonymisation decision making framework is a new approach to anonymisation that demystifies the complex but nevertheless tractable process of carrying out an anonymised release or share of data. The framework will continue to grow and develop; shortly we will be releasing statements about the relationship between the

framework and GDPR and we are also looking to produce editions contextualised to non-UK legislations.

References

ARRINGTON, M. (2006) *AOL proudly releases massive amounts of user search* Data TechCrunch, available at: http://tinyurl.com/AOL-SEARCH-BREACH [accessed 30/5/2016].

ATOKAR (2014) *Riding with the Stars: Passenger Privacy in the NYC Taxicab Dataset,* available at: http://tinyurl.com/NYC-TAXI-BREACH [accessed30/5/2016].

CAVOUKIAN, A. (2011). Privacy By Design: The 7 Foundational Principles, revised version, Toronto: Information and Privacy Commissioner of Ontario, https://www.ipc.on.ca/images/resources/7foundationalprinciples.pdf.

DUNCAN, G. & LAMBERT, D. (1989) The risk of disclosure for microdata; *Journal of Business & Economic Statistics*, 7(2): 207-217, DOI: 10.1080/07350015.1989.10509729.

EU GENERAL DATA PROTECTION REGULATION http://www.privacy-regulation.eu/en/index.htm

ELLIOT, M. J. & DALE, A. (1999) Scenarios of Attack: The Data Intruder's Perspective on Statistical Disclosure Risk; *Netherlands Official Statistics*, Spring 1999: 6-10, avilable at: http://tinyurl.com/ATTACK-SCENARIO [accessed 30/5/16].

ELLIOT, M.; LOMAX, S.; MACKEY, E. & PURDAM, K. (2010). "Data environment analysis and the key variable mapping system" in Josep Domingo-Ferrer & Emmanouil Magkos (eds.), Privacy in Statistical Databases, Berlin Heidelberg: Springer, 138-147.

ELLIOT, M.; MACKEY, E.; O'HARA, K. & TUDOR, C. (2016a). The Anonymisation Decision-Making Framework, Manchester: UKAN.

NARAYANAN, A. & SHMATIKOV, V. (2010). 'Myths and fallacies of "personally identifiable information"', Communications of the ACM, 53(6), 24-26.
OHM, P. (2010) Broken Promises of Privacy: Responding to the Surprising Failure of Anonymization; *UCLA Law Review*, 57: 1701-1777, available at: http://www.uclalawreview.org/pdf/57-6-3.pdf [accessed 30/5/2016]

PAASS, G. (1988). 'Disclosure risk and disclosure avoidance for microdata', Journal of Business and Economic Statistics, 6(4): 487-500.

REITER, J.P. (2005) Estimating Risks of Identification Disclosure in Microdata; *Journal of the American Statistical Association* 100(472): 1103-1112. DOI: 10.1198/016214505000000619.

UK: *Data Protection Act* (1998) London: The Stationery Office [Online], available at: http://tinyurl.com/UK-DPA98 [Accessed: 20/5/2016].

UK: INFORMATION COMMISSIONER'S OFFICE (2012a) *Anonymisation: managing data protection risk code of practice,* available at http://tinyurl.com/ICO-ANON [accessed 25/5/2016].