Metrics of Privacy Conference

14 November 2018

Report of a discussion meeting organised by the Technical Advisory Panel of the Investigatory Powers Commissioner's Office on Metrics of Privacy.

Participants:

Sir Bernard Silverman, Chair of IPCO TAP
Sir Adrian Fulford, IPC
Arosha Bandara, Open University
Paul Bernal, University of East Anglia
Mark Briers, Alan Turing Institute
Ian Brown, DCMS
Jonathan Cave, University of Warwick
John Davies, IPCO TAP
Julian Huppert, Director of Intellectual Forum
Eric King, Visiting Lecturer, Queen Mary University London
Derek McAuley, IPCO TAP
Javier Ruiz, Policy Director, Open Rights Group

Introduction

1. How can you measure intrusion into privacy? This was the question posed to those attending the Metrics of Privacy Conference held at Jesus College, Cambridge on 14 November 2018. The conference was organised by the Technology Advisory Panel (TAP), a statutory body set up in 2018 under the Investigatory Powers Act 2016[1]. Co-sponsorship was provided by the Intellectual Forum based at Jesus College, Horizon Digital Economy Research and the EPSRC IoT Research Hub, PETRAS. Delegates were invited from a range of backgrounds including statistics, computer science, and privacy. They brought experience both of academia, and of civil society groups and other bodies with an interest in the topic. The conference took place under the Chatham House rule, and this note summarises the discussions without necessarily implying that every participant agreed with every point recorded.

2. In recent years the subject of personal privacy has grown in importance and public focus, particularly in the light of WikiLeaks, the Edward Snowden revelations (2013) and the increased number of cases against the UK Government being taken to court by individuals and privacy organisations. This is alongside a background of the development of the new Investigatory Powers Act 2016 (IPA) and the growth of public scrutiny into the activities of the Intelligence Agencies and other organisations using surveillance powers; parallel to this is the need for the UK to protect itself against terrorism, cyber and other organised crime.

[1] The Technology Advisory Panel (TAP) is an independent group set up under the Investigatory Powers Act 2016 (paras 246-247).

3. Public expectation of privacy has been brought into question as never before. The IPA includes safeguards alongside each element of the powers, and the Investigatory Powers Commissioner (IPC) and his office (IPCO) are required to oversee how the powers are used and whether the safeguards are being correctly applied. The TAP, working alongside IPCO, has two functions: to give specific advice to the IPC on technical matters, and to take a broader view of technical development of relevance to the IPA. IPCO need to be able to measure, weigh, rank and quantify activities under the Act***; the provision of metrics to assess the level of privacy would be valuable*** and would aid IPCO both in correct assessment of warrant applications and in inspecting the use of the IPA powers. However, **there is no simple quantification method**.

4. Prior to the day, all participants had been asked to submit two questions for discussion via the Well Sorted tool[2]. The output from the tool produces a dendrogram which groups and sorts the submitted questions by frequency and topic of the questions.

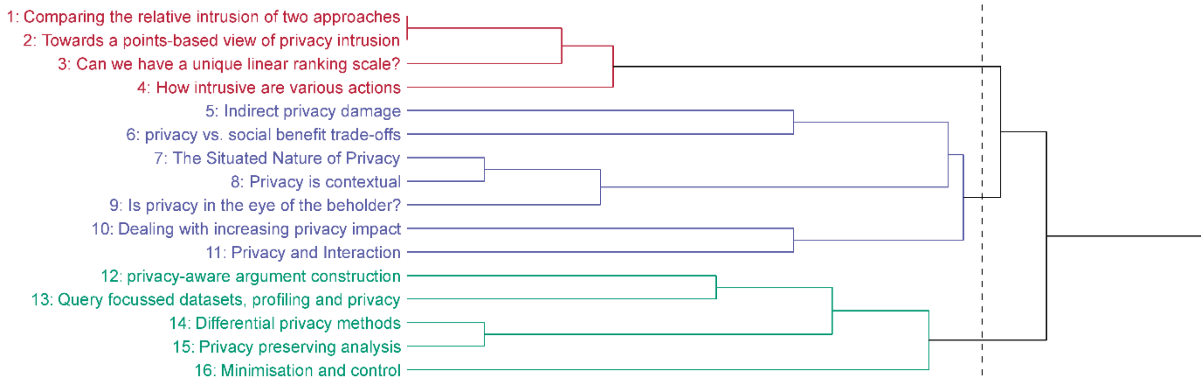
1: Comparing the relative intrusion of two approaches
2: Towards a points-based view of privacy intrusion
3: Can we have a unique linear ranking scale?
4: How intrusive are various actions
5: Indirect privacy damage
6: privacy vs. social benefit trade-offs
7: The Situated Nature of Privacy
8: Privacy is contextual
9: Is privacy in the eye of the beholder?
10: Dealing with increasing privacy impact
11: Privacy and Interaction
12: privacy-aware argument construction
13: Query focussed datasets, profiling and privacy
14: Differential privacy methods
15: Privacy preserving analysis
16: Minimisation and control

Three main themes were identified from this and taken as broad discussion themes for the day:

i. (In blue). **Privacy is Contextual**.
ii. (In green) **Mechanisms to mitigate privacy concerns**
iii. (In red). **Metrics**

Privacy is Contextual

5. The group was in full agreement that privacy is affected by the context. This is a broad topic and one on which people hold different views. Certain professions already have extra legal protections in place in recognition that greater privacy is essential to that context (e.g. journalists, lawyers and politicians). The topic is interdisciplinary and has a psychological angle and people's responses to privacy

[2] This is a tool devised by Heriot-Watt University as an aid to setting agendas for conferences which is widely used by the Research Councils

issues will vary based on their personality as well as the context. This makes the methodology of research into metrics much harder, as if privacy is mentioned, the response can alter. Additionally, the context in which responses are made may make comparisons (between them) difficult.

6. As the IP Act recognises, intrusion into privacy by the state is not always negative, so long as it is done in pursuit of legitimate aims and in accordance with the law. There are times when use of personal data may be of wider public benefit, even if it is intrusive. Privacy can therefore come at a cost. There was wide agreement that aiming never to have intrusive tools was not realistic, but that there should be an aim to minimise intrusion. The benefits may be perceived differently by different people, so the scale of privacy versus benefit needs to be very carefully balanced.

7. ***Intrusion is not simply an individual matter as it can involve not just the personal data of the individual but can have impacts on the wider group or population where there is some form of shared identity***. It can involve networks of individuals or groups or information networks. This element of wider impact and potential indirect effects needs to be considered in any attempt to define privacy metrics. An example of this might be the use of a DNA testing site such as 23andMe[3] where the addition of one's DNA may impact on others (such as highlighting familial links, identifying genetic traits or health-related issues which others may prefer not to know). Such issues might be managed by controlling any inferences from the DNA rather than by preventing the DNA itself from being put up on the site.

8. Where is the sensitivity? Is it in the data itself or in the analysis and the end result? It will be increasingly difficult with the growth of Artificial Intelligence (AI) to know what analytical work has been done on the data. The intrusion caused by obtaining and retaining the data is not a fixed impact but will vary according to the people whose data it is and what other data is available and may be combined with the original data. ***It is essential to reassess the potential for intrusion constantly*** as analytic processes change and develop. This can be more difficult when it involves large and therefore less flexible bureaucratic organisations such as Government. If the privacy intrusion is continuous this is equally important. There is a tension between privacy and accountability: it must be possible to be able to challenge a decision.

9. There are several layers to privacy; an individual may choose to reveal their personal data to others and this may affect the decisions of others. It is possible to get a great deal of knowledge about an individual and their behaviours whilst not knowing who they are; conversely, a few significant pieces of disparate data may

[3] 23andMe (www.23andme.com) markets direct-to-consumer genetic testing to consumers. The services offered include "health", "traits" and "ancestry".

completely identify an individual. Therefore, even for publicly held data, people have rights of privacy. They may have volunteered the information without understanding the implications of its use at the time, or the potential future use of the data. New techniques or methods may be developed over time which may deviate from the original premise for the use of the data.

10. Intrusion into privacy involves a wide spectrum: it is not as simple as zero versus full intrusion, nor is it a linear process. ***Machine collection is NOT zero intrusion as it still invades privacy.*** Awareness of the collection taking place may encourage individuals to change their behaviours, and actions do have consequences. Machines can create impacts which humans cannot and vice versa, such as: dealing with volumes, taking decisions, the ability to recognise "lies" or to deal with (or even introduce) bias, and machines may not be fully trusted by humans in how they manipulate the data. A machine can analyse more data more quickly, but in common with humans can make errors or miss things.

Differential Privacy and Mechanisms to mitigate privacy concerns

11. Is there a difference between the invasion of privacy done by a machine versus by a person? What if the data is merely processed and not viewed by a human? Is the processing of the data more intrusive than the collection? The nature of the intrusion is different[4] and each can be worse at different times. Several stages of processing might be employed, bringing increased interference with privacy at each stage. The political environment also may change how the data is used, which may be another argument for minimisation at the earliest possible stage. This prevents any risk of accidental or deliberate misuse.

12. How much is data deletion a mitigation of intrusion? Given the non-trivial cost of immediate destruction of data from large systems and traditional backup systems, to what extent does locking the data beyond use to be deleted by automated processes in the course of time (e.g. when backup tapes etc are recycled, or through set deletion timings) compare to immediate destruction? Is it possible to delete the data whilst retaining the results, and are these results more, or less, harmful that the data itself? This could make the data less revealing but would prevent later analysis or cross-linking of information. Time-delayed destruction can potentially be reversed, whereas destruction is final. The antithesis to this might be a case where the data if retained may prevent a miscarriage of justice due to incomplete or misconstrued earlier analysis. There may also be good intelligence reasons for retaining data, for national security purposes.

[4] Based on evidence to the Joint Committee on the IP Bill. See: Draft Investigatory Powers Bill Joint Committee, Volume of Written Evidence, 11th February 2016

13. A paper entitled “Technical Privacy Metrics”[5] was referenced during this session and recommended for further reading. In this document, eighty different measures of privacy are noted under eight separate categories. These are:
    a. Uncertainty
    b. Understanding the measures
    c. Data similarities
    d. Interaction of time
    e. Quantification of error
    f. Indistinguishability
    g. Accuracy and precision, and
    h. Success probability.

14. Obfuscation approaches, such as differential privacy[6] (e.g. adding random noise to data), could protect the individual whilst retaining the underlying population statistics but ideally need to avoid any temporal correlation ability. The approach of k-anonymity[7] might have value, allowing data to be processed to the point where you can address a group’s statistics without identifying any individual. However, reidentifying after anonymisation depends on what other data can be brought to bear, and there are no deeply theoretical results on such techniques on which to rely. A single piece of additional data may be enough to unlock data which has previously been anonymous. Encryption techniques could also be used to support privacy; however, metrics which appear secure now may not be in the longer term (e.g. encryption using 56-bit keys was once considered sufficient) therefore any metrics will need to be revisited constantly. Each of these could be applied in different contexts.

15. What is required is a process with real impact on expected intrusion and with the ability to oversee this effectively. How much should we be prepared to put trust in the agencies or indeed in the overseers? All are keen to ensure that intrusion is reduced as much as possible. There is always a risk of misuse and no perfect solution will be found; however, technical and architectural means can be used to make improvements. Some risks can be detected and stopped, others would need to be understood and properly managed by adapting the processes.

16. The focus has moved from cyber security to cyber resilience, balancing the risks against the utility, and providing metrics to help the decision-makers. The paper quoted above contains some metrics options which are quite simplistic, and simple metrics should be considered as much as more complex methodologies. If measures and metrics are extremely complicated, then decision-making can

[5] Technical Privacy Metrics: A Systematic Survey

[6] Differential privacy adds mathematical noise to a sample of the individual’s usage pattern. It aims to discover the usage patterns of a large number of users, without compromising individual privacy

[7] A dataset has the k-anonymity property if the minimum set size that shares the same characteristics is k; e.g. if personal data, each person cannot be distinguished from k-1 others.

become more difficult and highly disputable. Whatever metrics are proposed need to be transparent and to allow people to validate them by checking out their accuracy and effectiveness. Any metrics need to be deployable by the agencies and/or by the inspectors (either in oversight or else run for themselves) and must be resilient and not able to be subverted. We must build a correspondence of the risk thresholds already in place to the metrics, in essence defining a "privacy budget". However, ***such a correspondence would need to be regularly monitored so that it is consistently applied*** **as well as deal with improvements in the metrics themselves**. It is better to say that the measure of privacy will be improved, than to say that intrusion will be mitigated.

17. To consider specific metrics it might be helpful to concentrate on a specific capability or technique, as there is no universal metric. Is it appropriate to publicise which privacy metrics are being used? Is it possible to look at information in the public domain to consider potential privacy metrics? The IPC is in a strong position to require metrics and privacy to be included on agendas and plans for new systems. There is no limit to what questions IPCO can ask in terms of necessity and proportionality. However, engineering staff may not be the best people to consider what privacy measures could be installed therefore guidance from other bodies would be helpful, particularly where there are cross-system impacts. ***How can architecture be used to ensure that privacy is protected?***
    a. ***What is being done to protect privacy?***
    b. ***What consideration and balance has been put into the system design?***
    c. ***How can we measure privacy versus benefit?***
    d. ***What consideration is given to whether this is the best method to measure privacy (technical/policy /legal)?***

18. To obfuscate which metrics are in place by the agencies might it be possible to publish a range of possible options and methods? How much is it possible to compare and discuss methods meaningfully with similar organisations and other groups working on privacy issues across the globe? ***What research is being done based on publicly available data?*** Is effectiveness more important than privacy? How can measures be publicly validated, particularly where techniques are sensitive? The goal may be not to expose the privacy measures but to show improvement year on year. Trying out techniques is an engineering problem but testing out the techniques is a metrics problem which needs to use actual data. IPCO has a role in helping to define privacy intrusion and TAP has a role in helping the agencies to come up with varied methods and to be able to compare the intrusion level of each and prove that compliance standards are improving.

19. There is a need to build public trust and that has many more steps than any technical solution. It requires the input of a wide range of people besides engineers and certainly needs ongoing public discussions. ***How do we engage more (and***

more active) public discussion on this topic? This must include not just general discussion but more technical involvement.

<u>Metrics</u>

20. Can we compare two methods to see which results in less intrusion? Can we use global statistics, or do we need to take more personal issues into account? Should measurements differentiate across society? Global population statistics may not be the right answer: we need to understand and nuance the data; correlations may have enormous consequences e.g. it is possible to identify an individual from four separate data location points. This may be innocuous to some but of concern to others.

21. It would be sensible to consider incremental processing (if we do a, then b, then c…). For intelligence minimisation it is possible to take a lot of data which can show a range of different outcomes, or to restrict the data so that it cannot be used for certain purposes. As a democratic society, constraints need to be included and what is possible is constantly changing. The more the data is touched, the greater the interference. However, to reach the intelligence goal it may be necessary to evaluate a whole range of data first.

22. There was a suggestion of the possible use of "citizen juries" to look at case studies (though there might need to be caveats from the agencies).

23. A starting point for future discussion might be to take the following thoughts/topics for all to consider:

    a. How is data traditionally minimised?
    b. How is data selection and filtering done?
    c. Indexing
    d. Data deletion processes
    e. Datamining

    This needs to include ideas on correlations between algorithms and consideration of the risk of bias. Are there any non-technical means to reduce the level of intrusion?

24. It might be helpful to put together a list of valid questions which JCs/inspectors may wish to ask themselves. Can this group agree some core principles for this and provide a guidance document which also devises means for how the responses can be measured?

<u>Next steps</u>

25. It was agreed: to produce and agree this report of the discussion meeting; that participants would be willing to engage in further meetings to illuminate more detail, but it was agreed that these events would need to focus in on specific topics per meeting; and IPCO TAP agreed to consider how to proceed.

Conclusion

26. In conclusion, the discussions illuminated the inherent complexity in determining Metrics of privacy. They will be neither straightforward nor easy to implement, and there are a large number of factors to be taken into account. However, this was a useful and interesting workshop which has captured some of the complexity of the topic and given helpful pointers for future work.

Post meeting comment

27. A further thought received post meeting was how could investigatory patterns be represented (e.g.[8]) in a manner that would allow a methodology for privacy impact measures to be applied to these patterns. This would enable a “privacy by design” approach to specifying investigatory needs[9]. Additionally, the measurement approaches should allow privacy risks to be balanced against the social benefits (e.g.[10]) associated with the investigatory request. This approach could facilitate sharing of investigatory patters with a wider community of stakeholders to validate the privacy impact measures.

[8] Alrimawi, Faeq; Pasquale, Liliana and Nuseibeh, Bashar (2017). Software Engineering Challenges for Investigating Cyber-Physical Incidents.

[9] George Danezis, Josep Domingo-Ferrer, Marit Hansen, Jaap-Henk Hoepman, Daniel Le Métayer, Rodica Tirtea, Stefan Schiffner, Privacy and Data Protection by Design.

[10] Yang, Mu; Yu, Yijun; Bandara, Arosha and Nuseibeh, Bashar (2014) Adaptive sharing for online social networks: a trade-off between privacy risk and social benefit.

Bibliography

Some literature/reference materials relating to these ideas include:

- Alrimawi, Faeq; Pasquale, Liliana and Nuseibeh, Bashar (2017). Software Engineering Challenges for Investigating Cyber-Physical Incidents. In: *Proceedings of the 3rd International Workshop on Software Engineering for Smart Cyber-Physical Systems (SEsCPS '17)*, 20-28 May 2017, Buenos Aires, Argentina.
- Bernal, Paul: The Internet, Warts and All: Free Speech, Privacy and Truth, Cambridge University Press, 2018.
- Benford,S; Rodden, T; Calder, M; Sevegnani, M: On lions, impala, and bigraphs: modelling interactions in physical/virtual spaces.*ACM Transactions on Computer-Human Interaction, vol. 23, issue 2,* May 2016.
- Bradford Franklin, Sharon and King, Eric: Strategies for Engagement Between Civil Society and Intelligence Oversight Bodies, *New America*, last updated November 2018. *[Context of wider engagement for oversight]*
- Danezis, George; Domingo-Ferrer, Josep; Hansen, Marit; Hoepman, Jaap-Henk; Le Métayer, Daniel; Tirtea, Rodica and Schiffner, Stefan Privacy and Data Protection by Design, *ENISA Report*, December 2014.
- Draft Investigatory Powers Bill Joint Committee, Volume of Written Evidence, 11th February 2016
- National Research Council. 2015. Bulk Collection of Signals Intelligence: Technical Options. Washington, DC: The National Academies Press. *[The major US government-commissioned paper exploring what other methods could be used to reduce privacy intrusion at the point of collection].*
- Wagner, Isabel and Eckhoff, David: Technical Privacy Metrics: A Systematic Survey, *ACM Computing Surveys Volume 51 Issue 3*, July 2018 DOI 10.1145/316838
- Wetzling, Thorsten and Vieth, Kilian: Upping the Ante on Bulk Surveillance: An International Compendium of Good Legal Safeguards and Oversight Innovations, *Heinrich Böll Stiftung*, November 08, 2018 *[Non-technical methods that could reduce intrusion]*
- Yang, Mu; Yu, Yijun; Bandara, Arosha and Nuseibeh, Bashar (2014). Adaptive sharing for online social networks: a trade-off between privacy risk and social benefit. In: *13th IEEE International Conference on Trust, Security and Privacy in Computing and Communications (IEEE TrustCom-14)*, 24-26 Sep 2014, Beijing, China, IEEE, pp. 45–52.