

Decomposing Data Privacy for Evaluation

Karen Sollins

MIT Computer Science and Artificial Intelligence Laboratory

October 3, 2016

Extended Abstract

Submitted to the Federal Trade Commission for consideration at PrivacyCon 2017

Abstract:

One of the challenges we face in the provision of data privacy is the question of how effective we are at doing that. To tackle that problem in this paper, we propose a decomposition of the data privacy problem in three ways. The first is to identify the stakeholders in both the definition and the provision of privacy. The second is to consider the data life-cycle stages, because it is at those points that possible comparable approaches will be applied. The third is the identification of contexts that constrain or define privacy policies. In this decompositional framing of privacy, we examine approaches to the provision of privacy and ask about the metrics for evaluation of their efficacy and costs.

I. Introduction

The issue we address in this paper is how to evaluate the effectiveness, utility, and costs of providing data privacy in large scale, distributed, multi-function systems. There are three primary reasons for evaluating such privacy: social expectations, regulation and other governmental forces, and economics. There are social norms and expectations that will expect norms of privacy preservation. Governments may require privacy either in the form of regulation of other organizations or as part of their own missions. Finally, there may be strong economic and business drivers for privacy. In all three situations, being able to trust the efficacy and understand the

tradeoffs in providing privacy lead to a need for evaluation of the policies and mechanisms defining data privacy.

To do this, we must tease apart several key, more specific questions. In this work, we begin by examining the question of who cares, in whose interest is some form of privacy being provided and in what ways do they care. This is a question of the stakeholders involved in the provision and decision-making with respect to privacy. That is followed by framing the points at which privacy provision and possibly violation will occur. In this work, we use the data life-cycle to identify the points at which privacy “events” can occur. In terms of teasing apart the high level question of who cares about what, our third component is that of context, in the sense that Nissenbaum [3, 4, 5] uses the term. This is a means to distinguishing among privacy policies based on distinct policy domains, in addition to our previous two criteria.

With the three dimensional taxonomy of privacy laid out, we then turn our attention to the questions of approaches to evaluation. We identify and discuss three, with our primary focus on the third. These are (1) careful examination and approximate analysis; (2) proofs of logical correctness; and (3) metrics and measurement. The remainder of the work focuses on these metrics and measurements, with an examination of the different kinds of metrics that may be applicable to different elements of the taxonomy.

Although we are a long way from being able to do this, we are driving toward being able to understand the composition of these kinds of evaluation, in order to gain an overarching answer to the question of how private is the composition of a set of privacy preserving techniques. We definitely must leave that to future work, and focus here on the decomposition questions with respect to privacy.

II. Decomposing privacy in three dimensions

The problem domain we are working in is the provision of policies with respect to information or data that give the stakeholders in the data confidence in an acceptable model of privacy. In order to be able to evaluate the effectiveness of provision of privacy, we must analyze who cares about the privacy, what are the privacy risks to the data, and how do those vary depending on the situation or context. Thus, we reflect on the roles of the stakeholders involved, the vulnerabilities of the data at different stages, and finally, the utility of defining distinct contexts across both stakeholders and stages in the handling of the data.

A. Who cares? The Stakeholders

Privacy concerns will vary depending on the role that an individual has with respect to the data. Thus, for example, if the data is data about an individual only, then that individual probably has primary interest in privacy. That said, if the primary individual is a child, then his or her parents may both have significant interest and may find that secondarily they are the subjects of the data. In an earlier report, [1], we examined a series of privacy scenarios in the context of big data. In [2], we refined the list a bit, but it derives from the use scenarios. Although we recognize that not all data is “big data”, those scenarios enabled us to distinguish a set of qualitatively different stakeholder roles in the privacy domain. That list provides the core of

our stakeholder analysis. It is important to recognize that some stakeholder roles may not exist in some circumstances. In others roles may be merged. Our list of stakeholders includes:

- Data subject(s): the people represented in the data; they are sometimes the primary source of policy statements, but may also include secondary subjects;
- Decision-makers: those making the determination to collect the data and are often the “data beneficiaries”; they also often are a significant source of policies;
- Data regulators: those with a role in determining data collection, sharing, and use policies broadly;
- Data collectors: those collecting the data; these may play this role by choice, or as required by regulation;
- Data curators: those determining the correctness, completeness and other validity checks of the data;
- Data analysts: those using the data for learning, inference, and other analysis;
- Data platform providers: those storing and managing the data, running the infrastructure for the data;
- Policy enforcers: those determining what mechanisms are used to provide confidence in the application of the policies;
- Auditors: those who evaluate whether or not policies have been applied appropriately, with respect to the policies that are applicable.

B. When are policies applied? The data life-cycle

Our second dimension to decomposing privacy is to identify the stages of the data life-cycle. We originally chose this decomposition in order to categorize the applicability and utility of different technologies being proposed to “provide privacy”. [2] We identify six stages in the data life-cycle and then for each begin to identify categories technologies to provide some elements of privacy at each.

- Data collection: online notice and consent, informed consent;
- Data management: system design, encryption technologies and other system level privacy enforcement;
- Data access: Data use agreements, authentication and authorization protocols, hardware authentication such as USB keys and biometrics, standard encryption techniques;
- Data processing: If the intended use is insights about an aggregate population only, various statistical approaches such as Differential privacy and synthetic data sets (e.g. Bayesian statistics). If the intended use is both about aggregate populations and individuals, then personal and private data stores, secure multi-party computation and functional and homomorphic encryption are possible technologies;
- Data compliance and audit: User access logging, automated policy analysis and applicability. An example of this is the compliance bootstrapping in [6];
- Data destruction: guarantees of deletion, destruction of encryption keys.

Again, as with the stakeholders, in some cases, some of these may be merged or collapsed.

C. How does the situation contribute: Context

The third element of our privacy evaluation framework is the concept of context as derived from Nissenbaum's work on *contextual integrity* [3, 4, 5], as well as our own more generalized concept of a *region*, intended as a first class abstraction in managing networks generally. In both of these works, the objective is to define localized explicit scopes of applicability. With respect to privacy, such a contextual or regionalized model reflects both the definition and application of privacy policies. As Nissenbaum amply exemplifies the same action by the same individuals may fall under and be bound by significantly different privacy policies depending on the context. This means that the evaluation of the

application of privacy policies will also be dependent on context. Being able to identify those contexts will be an important component of the evaluation of the efficacy of privacy policies. Thus, as Barth et al. [8] do, a formalism will play a significant role here. Attention to that is left to a more complete version of this work.

D. Approaches to evaluation

Overall, in our discipline, we can identify three primary approaches to evaluation, careful examination and approximate analysis, proofs of logical correctness, and metrics and measurement. Each is important and has its own value.

- Careful examination and approximate analysis is the approach of arguing by example and reasoning about the correctness and/or validity of an approach or design. We see this both broadly across the field of computer science and in much more detail, for example in the analysis of security vulnerabilities. See any number of IETF RFC protocol specifications security sections for examples of this approach.
- Proofs of correctness are often quite compelling, although often dependent on one or another specialized "logics". They provide a degree of confidence in logical reasoning about the efficacy or accuracy of what is being questioned, but have the drawback that they often tell the audience little about actual implementability or performance. The work of Sen et al. [6] is a concrete example of this as applied to privacy.
- Metrics and measurement is a third approach to evaluation. In this category of evaluation, one must determine what matters to whom, and probably especially, in what context. One then must determine how to actually measure the metric. This may take several forms ranging from simulations of possibly abstracted forms of the target, to active measurements in testbeds or other realistic situations to passive measurements in the wild. All

have their value. It is this third area where we focus the majority of this work here.

With this in mind, we consider what is being done at present in terms of metrics and measurements. This will allow us to ask questions about whether those are the right questions to ask, is there agreement on them, and whether there are gaps.

III. Evaluation of privacy through metrics and measurement

In this section, we review particular approaches and technologies proposed for providing privacy. Each tackles some part of the overall privacy challenge. Different stakeholders may play a role in one or another. Each can be placed in one part of the data life-cycle, and generally each is best evaluated within particular contexts. Each has distinct metrics used for evaluating its effectiveness. For purposes of this extended abstract we highlight only three such examples, the work of Cranor on notice and consent, the work of Sen et al. mentioned earlier on the development of the *Legalease/Grok* toolkit, and the work of Dwork on differential privacy. The complete version of this work includes technologies at all stages in the data life-cycle with selections from different contexts.

A. Defining and presenting privacy policies

The data life-cycle begins with decisions about the collection of data. This is done in the context of some privacy policy, which often, although not always is presented to the subject in some form. One of the leaders in evaluating the effectiveness of such policy presentation is Cranor, who has led a number of studies including [9, 10, 11, 12]. The kinds of metrics considered in this body of work focus on whether the user or customer either

understands or makes choices based on privacy policies. Therefore the kinds of metrics that are relevant are the impact of iconography vs. text (presence or absence of icons), related tagging, the timing of presentation (such as in an app store vs. inside the application itself), among others. The studies are done in contexts where the customer is doing something else, for example purchasing something or playing a game. The key is that the focus in this body of work is on evaluating either the customer's understanding of the privacy message or a reflection of customer decisions based on privacy as one factor among several.

There are two key points to make in this space. First, as Nissenbaum points out [5], there remains a significant challenge in evaluating the effectiveness on customers of the presentation of policies: there is a significant tradeoff with respect to transparency of policies. The dilemma is that the full detail of a policy may be so detailed and specific that one cannot effectively present it to customers. In contrast, what one can present to customers, especially if it is in the form of icons or "nutritional labels" has lost so much detail that the customer is not possibly informed of the critical privacy risks.

The second is that there are other stakeholders besides the customers for whom evaluation of the effectiveness of presentation of privacy policies may be important. These may include regulators and other decision-makers about the contextual definition of a privacy policy. Other evaluation metrics will be important for a more complete privacy evaluation, even of only the issue of the efficacy of notice and consent.

B. Evaluating compliance

There are at least two significantly different approaches to take in evaluating or enforcing compliance with respect to privacy policies, either prior to execution of the code and one ex post facto. In the first area, one can ask the question of whether certain conditions can be guaranteed to be met with respect to privacy

policies and how much it costs to do that. In the latter, accountability is central to tracking what behaviors occurred and whether they violated the policies; only with that can one provide recourse for policy violation. For purposes of this extended abstract we will focus only on an example of the first. Several key pieces of work on the second include Weitzner et al. [13] and more recently Datta et al. [14]. Weitzner et al. focuses on how to keep track of who has had what interaction with the data. Datta et al. focuses tracking the data, that is used in making a decision. Thus, if a subject is denied a mortgage, the second considers how to understand which data was used in that decision, with an eye toward avoiding discrimination or use of data that is otherwise considered private.

We focus here on the toolkit presented in Sen et al. [6]. That work addresses the problem of large-scale code compliance to privacy policies, in a MapReduce data processing model. It presents two tools, *Legalease*, which encodes English privacy policies, often produced by lawyers, into a policy expression language. The second tool *Grok*, in turn evaluates whether the output of *Legalease* is correctly embodied in the code of a data center, by mapping the data flow in code level data types into the abstractions from the *Legalease* representation, in order to evaluate conformance of the data flow to the policies.

The kinds of metrics for evaluation discussed in this work range from lines of code and nodes in the data structures generated to usability (tested on a set of security expert, but novice user volunteers) to expressiveness to precision. In this work, conservatism was given preference of precision. In addition, the precision required by the *Legalease* component of the approach leads to a minimization of false negatives.

C. Data processing and analytical methods

At the point in the data life-cycle when it is being processed and used, we begin by observing that if the intended use is at least in part to develop insights about individuals

then different tools and constraints will be applicable than if the intended use is only for large scale aggregate insights. In particular, as an example, the work originated by Dwork [15, 16]. The objective of this work is to provide a statistical probability ϵ that given two queries of a data set, one in which a record about an individual is present and the other in which that record is not present, with a tunable probability there is no observable difference. If that is true then with that level of probability, nothing about the individual is knowable from such a set of queries. In addition, the objective is that the dataset remain encrypted, so there is no chance of additional leakage of private information. In this work, the most interesting metric is the tunable ϵ , which is the basis for a set of tradeoffs between amount of work and degree of privacy. An interesting point here is that although ϵ is grounded beautifully in theory, it is in fact difficult to understand and how to choose it in any particular situation remains challenging. Furthermore, for the most part there are some limitations on the situations in which differential privacy is useful or effective. In general, it is dependent on a static data set, so using it with streaming data is problematic. In addition, there are questions about both the size of the dataset (must be large, by some measure) and the size of the query set (must be small by some measure). Finally, related to all this is how the choice of ϵ relates to *context* as we discussed early in this work.

D. The broader set of privacy components and their evaluation

For purposes of this extended abstract we have only highlighted a small subset of the full range of the elements of privacy, the capabilities used to address or support them, and approaches for evaluation of those. The full paper will contain a more extensive review of this space with a clearer delineation of the state of the art, where gaps may lie, where there are opportunities for further work, and where there are promising directions.

IV. What to conclude

We began this work with the objective of understanding how effective our online privacy is. The reasons for that are three-fold, social, regulatory and economic. It is well understood that privacy, whatever we mean by it, is unlikely to be perfect, and if it were, we would be unlikely to do many of the activities we enjoy or expect on a daily basis. At the core of providing some degree and model of privacy is understanding the objectives of that privacy. The social reasons extend to defining our friends and other social groups as boundaries for exposing information about ourselves. The regulatory reasons extend to the responsibilities of governmental regulators for different kinds of data about individuals and how that can and cannot be used, in this case with respect to provision of or violation of privacy. Thus, the government determines different privacy constraints for financial and healthcare services, as well as for the activities of the government itself in doing its jobs. These may include taxation, intelligence, security, and a large number of other governmental responsibilities. Finally, one must consider the economics of privacy. One must ask about the economic tradeoffs, the costs and benefits. We have not even touched on that in this review. In each of these arenas, there may be number of different sorts of Nissenbaum type contexts, with different criteria and objectives for provision of privacy.

This work takes a first step at framing privacy in such a way that we can begin to examine how to evaluate privacy. We have taken a three-pronged approach, identifying a set of stakeholders, a set of stages in the life-cycle of data, and the concept of context, as three orthogonal components of the question of evaluating privacy.

References

1. Bruce, E., Sollins, K. (ed.), Vernon, M., Weitzner, D., **Big Data Privacy Scenarios**, MIT Big Data Initiative, MIT-CSAIL TR-2015-030, Oct., 2015.

2. Sollins, K., Bruce, E. **Privacy and Big Data: Review of Emerging Technologies**, under review
3. Nissenbaum, H. **Privacy as Contextual Integrity**, *Washington Law Review*, 2014.
4. Nissenbaum, H., **Privacy in Context: Technology, Policy and the Integrity of Social Life**, Stanford University Press, 2010)
5. Nissenbaum, H. **A Contextual Approach to Privacy Online**, *Daedalus* 140(4), Fall 2011.
6. Sen, S., Guha, S., Datta, A. Rajamani, S., Tsai, J., Wing, J., **Bootstrapping Privacy Compliance in Big Data Systems**, Proc. IEEE 35th Symposium on Security and Privacy, May 2014
7. Sollins, K., **An Architecture for Network Management**, Proc. ACM CoNext ReArch Workshop, Rome, Italy, 2009.
8. Barth, A., Datta, A., Mitchell, J., Nissenbaum, H., **Privacy and Contextual Integrity: Framework and Applications**, Proc. IEEE Symp. Security and Privacy, 2006.
9. Kelley, P., Cesca, L., Bresee, J., Cranor, L., **Standardizing Privacy Notices: An Online Study of the Nutrition Label Approach**. CHI2010. [Originally published as Carnegie Mellon CyLab Technical Report CMU-CyLab-09-014, November 10, 2009.]
10. Egelman, S., Tsai, J., Cranor, L., Acquisti, A., **Timing Is Everything? The Effects of Timing and Placement of Online Privacy Indicators**. CHI '09: Proceedings of the SIGCHI conference on Human Factors in Computing Systems, 2009.
11. Kelley, P., Cranor, L., Sadeh, N., **Privacy as Part of the App Decision-Making Process**. CHI 2013.
12. Balebako, R., Shay, R., Cranor, L., **Is Your Inseam a Biometric? A Case Study on the Role of Usability Studies in Developing Public Policy**. Workshop on Usable Security (USEC 2014). San Diego, CA, February 23, 2014.
13. Weitzner, D. Abelson, H. Berners-Lee, T. Feigenbaum, J., Hendler J. Sussman, G.,

Information Accountability, CACM
51(6), June 2008.

14. Datta, A., Sen, S., Zick, Y., **Algorithmic Transparency via Quantitative Input Influence**, in *Proceedings of 37th IEEE Symposium on Security and Privacy*, May 2016.
15. Dwork, C., **The Promise of Differential Privacy**, Proc. IEEE FOCS 2011.
16. Dwork, C., Roth, A., **The Algorithmic Foundations of Differential Privacy**, Foundations and Trends in Theoretical Computer Science, 9(3-4), 2014.