A Classification for Privacy Techniques

Carlisle Adams*

THIS PAPER PROPOSES A CLASSIFICATION for techniques that encourage, preserve, or enhance privacy in online environments. This classification encompasses both automated mechanisms (those that exclusively or primarily use computers and software to implement privacy techniques) and non-automated mechanisms (those that exclusively or primarily use human means to implement privacy techniques). We give examples of various techniques and show where they fit within this classification. The importance of such a classification is discussed along with its use as a tool for the comparison and evaluation of privacy techniques.

CET ARTICLE PROPOSE UNE CLASSIFICATION des techniques qui cherchent à encourager, à préserver et à améliorer la protection de la vie privée dans les environnements en ligne. Cette classification comprend des mécanismes à la fois automatisés (dont la mise en œuvre se fait exclusivement ou principalement à l'aide d'ordinateurs et de logiciels) et non automatisés (dont mise en œuvre se fait exclusivement ou principalement par l'intermédiaire de personnes). Des exemples sont donnés de diverses techniques, en les situant dans cette classification. L'article commente l'importance des classifications de ce genre ainsi que leur utilité pour la comparaison et l'évaluation des techniques pour la protection de la vie privée.

* School of Information Technology and Engineering (SITE), University of Ottawa, Canada. The author gratefully acknowledges the research assistance of Maxime Laverdière, now a law student at McGill University in Montreal, who provided many sanity checks and much fruitful discussion on this work while he was on contract to the Office of the Privacy Commissioner of Canada during the summer of 2004.

A Classification for Privacy Techniques

Carlisle Adams

1. INTRODUCTION

AS THE INTERNET AND THE WORLD WIDE WEB continue to expand, the average user has available to her a combination of processing power and data access that was unimaginable a relatively small number of years ago. Either of these in isolation represents a security danger in the hands of a user who is malicious. Significant processing power can allow key and password spaces to be searched easily, but only if encrypted text or the hashed password is readily available; significant data access can allow confidential information to be revealed, but only if sufficient compute cycles can be spent to find it. The combination, however, leads to serious danger not only to security, but also to privacy.

The average user, therefore, can, with relatively little effort, harness the tremendous processing power of the Web (particularly when computing nodes are linked in some way, such as by using a Grid network[1]) to analyze the vast repository of data available on the Web. Any personal information about a user Alice that is accessible by users somewhere on the Web can be found, analysed, modified, deleted, combined with other information, or disseminated more widely without Alice's knowledge or consent, potentially resulting in considerable damage or embarrassment to Alice. Because of this, it is widely recognized that the internet presents a serious threat to privacy throughout all segments of society.

In response to this threat, researchers in academia, government, private industry, and elsewhere have invested considerable effort in recent years toward devising and implementing techniques to protect privacy in online environments. Some of these techniques are software- or hardware-based, using algorithms and computers to preserve or enhance privacy. Other techniques are

1. See *e.g.*, Grid Computing Info Centre (GRID Infoware), <http://www.gridcomputing.com>.

human-based, using laws or guidelines to encourage good privacy practices in specific sectors, such as healthcare or finance. The amount of research activity in privacy has been growing steadily, almost exponentially,[2] over the past two decades. While in many respects this is very good news for this field, it has given rise to some difficulties. In particular, the highly multidisciplinary nature of privacy has made it hard for privacy advocates, lawyers, engineers, computer scientists, and others to discuss and compare ideas using common terminology and concepts. Furthermore, it can be difficult for any researcher to understand what relevance or significance a particular approach has to the overall goal of privacy and how it is related to other existing approaches in the field.

This paper attempts to make a step toward addressing these difficulties by proposing a high-level classification for privacy techniques. The goal of this classification is to organize these techniques in a manner that allows them to be more easily understood, compared, and analyzed. This classification encompasses both automated and non-automated techniques and can be used as a tool for explaining and evaluating a variety of privacy mechanisms.

*

2. DIFFICULTIES IN CREATING A CLASSIFICATION

THE TASK OF DEFINING A CLASSIFICATION for techniques that preserve privacy in online environments is not trivial. There are two main reasons for this (arising from the highly multidisciplinary nature of the privacy field):

1) sometimes there are many different views of the same concept across disciplines; and

2) sometimes there are fundamental differences between disciplines with respect to which concepts are important.

With respect to the former difficulty, the fact that specialists from widely different fields work on privacy inevitably leads to incompatible definitions of some concepts. While in most cases discordance or confusion can be cleared up with a simple discussion, at times there are fundamental dissimilarities that require substantial work and investigation to resolve. For instance, the key concept of what constitutes "personal information" has many different interpretations, some highly abstract and some too concrete and restrictive. In the *Privacy Act*,[3] for example, personal information is defined as "information about an identifiable individual that is recorded in any form."[4] However, this definition was later seen to be too restrictive; it led to unintended consequences

2. See Free Haven, "Anonymity Bibliography," <http://freehaven.net/anonbib/> [Haven, "Anonymity Bibliography"].
3. *Privacy Act*, R.S.C. 1985, c. P-21, <http://laws.justice.gc.ca/en/P-21/95414.html>.
4. *Ibid.*, s. 3.

in court because bodily fluids were not treated as personal information. The *Personal Information Protection and Electronic Documents Act*[5] (PIPEDA) therefore omits all reference to the word "recorded" and defines personal information as "information about an identifiable individual"[6] (a more general definition that perhaps encompasses too much, in that it includes "public" or readily-observable information about an individual that would not be regarded as "personal" in some environments). In technology circles,[7] personal information (or Personally Identifiable Information, PII) is defined as "any piece of information which can potentially be used to uniquely identify, contact, or locate a single person."[8] Note the shift in focus from information *about* a person to information *that can be used to pinpoint* a person. In particular, according to Wikipedia,[9] the following are not generally considered to be PII: first or last name, if common; country, state, or city of residence; age; gender; and race.[10] It is recognized, however, that combinations of such data may become identifying. Defining personal information in a way that is relatively precise and workable, and yet acceptable to multiple disciplines, is therefore a challenging task. The definition and notation I propose below is a result of close collaboration with engineers, computer scientists, managers, lawyers, and governmental policy makers; from our discussions thus far, it seems to strike the right balance between specificity and suitability for all parties.

The other main difficulty in creating a classification in this area is that universally defining the terminology is not sufficient to harmonize different views on privacy. In particular, some notions are considered critical and indispensable in one field, but have little or no relevance in another field. Technological mechanisms, for example, are centered on the idea of user control; a technology will be considered effective if it allows the data subject to maintain strict control over who has access to the data that this subject herself has determined to be sensitive. Thus, technological privacy mechanisms focus, for example, on allowing users to create their own privacy preferences, to choose specific recipients for data and encrypt that data for them alone, to establish anonymous channels for communication, and to determine whether another party's privacy practices are acceptable. In the legal domain, however, user control (in the form of consent or of active participation in the protection of the data) plays a relatively minor role. Consent is mentioned in only two ("Collection Limitation"[11]

---

5. *Personal Information Protection and Electronic Documents Act*, S.C. 2000, c. 5, <http://laws.justice.gc.ca/en/p-8.6/93196.html> [PIPEDA]. See also Stephanie Perrin *et al.*, *The Personal Information Protection and Electronic Documents Act: An Annotated Guide* (Toronto: Irwin Law, 2001).
6. PIPEDA, *supra* note 5, s. 2.
7. See generally, references to personal information in such sources as Wikipedia, <http://www.wikipedia.org> [Wikipedia].
8. Wikipedia, "Personally Identifiable Information," <http://en.wikipedia.org/wiki/Personally_identifying_information> [Wikipedia,"PII"].
9. Wikipedia, *supra* note 7.
10. Wikipedia, "PII", *supra* note 8.
11. *OECD Guidelines*, *infra* note 13 at p. 14.

and "Use Limitation"[12]) of the eight OECD[13] privacy principles, and in fact is not even necessary in some situations (in "Use Limitation," disclosure or use of personal data for purposes other than those originally specified may occur by authority of law without consent of the data subject[14]). Furthermore, the law may define what information is sensitive with respect to Alice, regardless of Alice's opinion on the matter. Finally, some laws (*e.g.*, HIPAA[15]) recognize accuracy, openness, and individual access as essential prerequisites to any privacy framework, whereas Privacy Enhancing Technologies (PETs) often neglect them.

This difference in emphasis between different domains brings to light many difficulties, not the least of which is the choice of discriminating factors in the classification. While a classification based on functionality may be a reasonable choice for technological methods, it is not very appropriate for legal methods, for example. Alternatively, a classification based on the entity responsible for initiating or enforcing a technique, which may be suitable for legal methods, is not ideal for technological methods, since often multiple entities must co-operate to provide privacy. The discriminators given in the proposed classification below appear to provide a useful classification across domains.

*

3. RELATED WORK

A SIGNIFICANT AMOUNT OF WORK has been done on proposing, improving, and implementing various PETs over the past twenty-five years.[16] This has led to a number of surveys, overviews, critical studies, and discussion papers, particularly in the last five to seven years.[17] However, very little work seems to have been done on categorizing these technologies in a systematic way for the purpose of comparison and analysis, and even less has been done on a categorization that includes other techniques addressing privacy, such as the legal infrastructure.

---

12. *Ibid.* at p. 15.
13. Organisation for Economic Co-operation and Development (OECD), OECD *Guidelines on the Protection of Privacy and Transborder Flows of Personal Data* (Paris: OECD Publications Service, 2001), <http://www1.oecd.org/publications/e-book/9302011E.PDF> at pp. 14–16 [*OECD Guidelines*].
14. *Ibid.* at p. 15.
15. *Health Insurance Portability and Accountability Act of 1996*, Pub. L. No. 104-191, Stat. 1936 (1996), <http://frwebgate.access.gpo.gov/cgi-bin/getdoc.cgi?dbname=104_cong_public_laws&docid=f:publ191.104.pdf> [HIPAA].
16. See Haven, "Anonymity Bibliography," *supra* note 2.
17. Philip E. Agre & Marc Rotenberg, eds., *Technology and Privacy: The New Landscape* (Cambridge: MIT Press, 1997); Ian Goldberg, David Wagner & Eric Brewer, "Privacy-Enhancing Technologies for the Internet" in *Proceedings of the 42nd IEEE International Computer Conference* (Washington, DC: IEEE Computer Society, 1997), <http://www.cs.berkeley.edu/~daw/papers/privacy-compcon97-www/privacy-html.html>; John Borking & Charles Raab, "Laws, PETs and Other Technologies for Privacy Protection" 2001:1 Journal of Information, Law & Technology <http://www2.warwick.ac.uk/fac/soc/law/elj/jilt/2001_1/borking/>; The EXOCOM Group Inc., "Privacy Technology Review," <http://www.hc-sc.gc.ca/ohih-bsi/pubs/2001_tech/tech_e.html>; Lorrie Faith Cranor, "The Role of Privacy Enhancing Technologies" in Paula J. Bruening, ed., *Considering Consumer Privacy: A Resource for Policymakers and Practitioners* (Washington, DC: Center for Democracy and Technology, 2003), <http://www.cdt.org/privacy/ccp/roleoftechnology1.pdf>.

One notable exception in this area is the work of Rezgui, Bouguettaya, and Eltoweissy.[18] In their recent paper, Rezgui *et al.*, propose a taxonomy of technology- and regulation-enabled solutions for privacy preservation in the Web.[19] Their approach categorizes solutions to the Web privacy problem based on the primary enablers of privacy preservation. As mentioned above, such a discriminator may be suitable for the legal branch of the taxonomy (where, for example, one can discuss "self-regulatory solutions" and "mandatory regulations solutions"), but is less useful for the technology branch. Rezgui *et al*, describe client-based solutions, server-based solutions, and client-server solutions on the technology side of their taxonomy, but there are at least two difficulties with such a classification. First, in much of the technology community, designating a node as a "client" implies that some other node will be a "server" and so it seems somewhat counter-intuitive to include peer-to-peer architectures (such as many of the anonymizer architectures) in which there are no clients or servers, just peer nodes, under "client-based solutions." Second, and more importantly, this classification has no further role for the *enabler* discriminator: technologies are listed without further sub-classification under client-based, server-based, and client-server. Thus, it is difficult to know whether this taxonomy is complete and it is virtually impossible to compare adjacent technologies in any meaningful way. For example, how can one usefully compare personal firewalls and remailers, or P3P and PGP?

The classification proposed below is based on a series of discriminators that, ultimately, allow technologies within a category and between categories to be compared and contrasted in meaningful ways. It thus overcomes the difficulties encountered with the taxonomy of Rezgui *et al*, and provides an intuitive categorization of privacy techniques for online environments.

*

4. PERSONAL INFORMATION AND PRIVACY

I BEGIN WITH A DEFINITION of personal information. PIPEDA,[20] a Canadian law focusing on the privacy of personal information, defines personal information as *information about an identifiable individual*.[21] While this definition is suitable for many purposes, it is not sufficiently general in some cases. I therefore clarify and generalize the concept of an "identifiable individual." First, information about a group may be considered to be personal information with respect to a member of that group. For example, the statement "the Jones family is bankrupt" is likely to be regarded as personal information by Alice Jones if she is known to be a member of the Jones family referred to in the statement. Thus, personal information may be about several entities simultaneously without naming any of the individual entities explicitly. Although this may be implied in the notion of

18. Abdelmounaam Rezgui, Athman Bouguettaya & Mohamed Y. Eltoweissy, "Privacy on the Web: Facts, Challenges, and Solutions" (2003) 1:6 EEE Security and Privacy 40.

19. *Ibid.*

20. PIPEDA, *supra* note 5.

21. *Ibid.*, s. 2.

"identifiable individual," I make this more precise in the definition given below. Second, there are situations in which an entity may treat some information as her personal information even though she is not the identifiable individual who is the subject of the information (see below). The following definition also deals with such cases.

Let E be a set of entities (more precisely, the set of identities associated with a set of entities), where $|E| \geq 1$. Let A be the identity of Alice and let B be the identity of Bob. Furthermore, either $A \in E$, or $B \in E$ and Alice $\in \mathfrak{R}$(Bob); that is, either Alice's identity is in the set E, or Bob's identity is in the set E and Alice is a member of the set of valid *representatives* of Bob (for example, Bob is a minor and Alice is his legal parent or guardian, or Bob is unfit or incapacitated in some way and Alice is his legal power-of-attorney). Finally, let I be some information that contains or implies the set E along with some other data that is associated with E. The information I is then *personal information*. Note that if $A \in E$ then I may be referred to as "Alice's personal information," whereas if Alice $\in \mathfrak{R}$(Bob) and $B \in E$ then I is not technically Alice's personal information, but is really "Bob's personal information" (over which Alice has valid legal authority). In some environments the distinction is not critical (in the case of a minor and a legal parent/guardian, for example, the only legally-recognized authoritative voice is that of the parent/guardian and so, for legal purposes, there is no distinction); in other environments the distinction may be more important (for example, a manager or agent that represents a client in the entertainment or professional sports industries may have some authority over the client's personal data, but the client will retain ultimate authority over this data). For the purposes of this paper, however, the distinction is ignored and I is loosely referred to in both cases as "Alice's personal information."

Given the existence of I, there is a group of entities (typically Alice, along with lawmakers and other entities at the regional, national, and international government level) that explicitly or implicitly defines one or more sets of valid recipients, R_j, for I and a set of valid purposes, P_{Rj}, for which I may be used by the recipients in R_j, for each j. *Privacy*, then, can be understood with respect to the above definition of personal information. Let r be a receiver of I (*i.e.*, r has acquired this information by some means) who uses I for some purpose, p. A *breach of privacy* has occurred if and only if $r \notin R_j$ for any j, or $p \notin P_{Rj}$ when $r \in R_j$ for some j (that is, r is not a valid recipient, or r uses I for a purpose that is invalid for his/her recipient set). Avoiding breaches of privacy is the process of exercising control over who receives personal information and how it is used. A *privacy technique* is a mechanism (that may be employed by Alice or by others) to enable such control; that is, it is a mechanism for restricting the recipients of I and the purposes for which I may be used to defined sets R_j and P_{Rj}.

*

5. CLASSIFICATION

THE CLASSIFICATION PROPOSED in this paper is shown in Figure 1. This classification is for techniques that encourage, preserve, or enhance privacy in online environments. Thus, activities such as wearing dark sunglasses and a false moustache in public or using cash for purchases at a convenience store, although they are both privacy-preserving techniques, are intentionally outside the scope of this work. Online environments support a variety of activities including electronic communications such as email and other forms of messaging, electronic shopping and auctions, electronic banking and finance, electronic delivery of entertainment and games, electronic learning and education, electronic healthcare, and the use of Web portals and search engines. In all such environments there can be a requirement for privacy, and a wide variety of techniques has been proposed over the years to address this need. As stated above, the goal of the classification proposed here is to organize these techniques in a manner that allows them to be more easily understood, compared, and analysed.

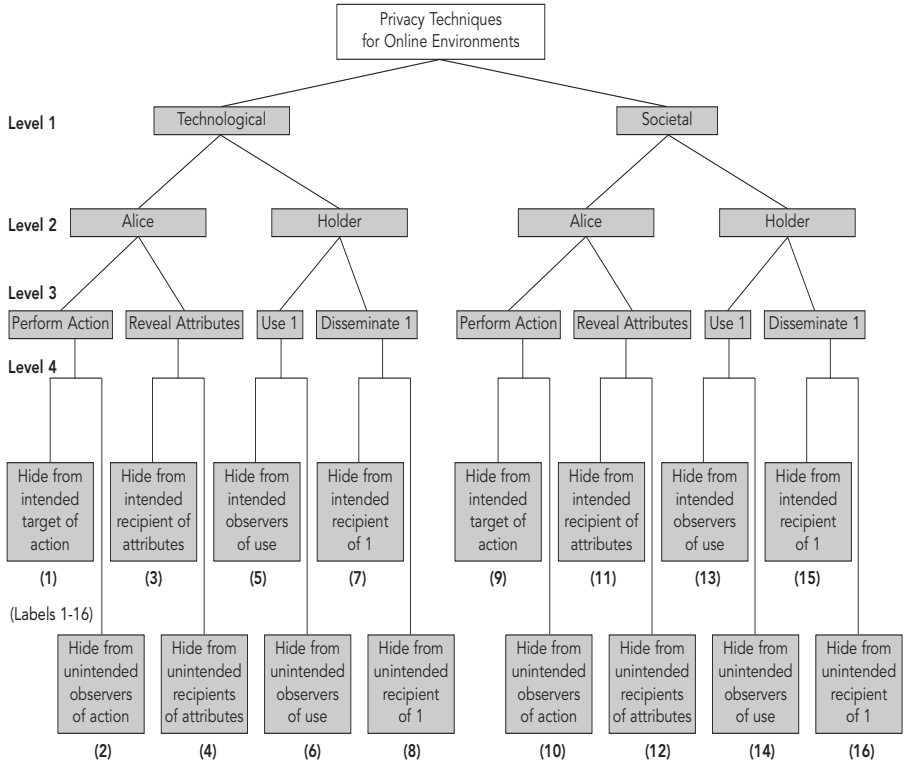

Figure 1. Classification of privacy techniques for online environments

I begin with a discussion of the intuition behind each of the levels in the classification. The following subsection then presents examples of actual privacy techniques and how they fit within this classification.

Level 1. The first division in the classification is between *technological* privacy techniques and *societal* privacy techniques. A technique is classified as technological in nature if it is enabled primarily or exclusively by machines (software or hardware running on computing equipment) with minimal intervention by humans once an initial installation and set-up phase has been completed. On the other hand, a technique is classified as societal in nature if it is enabled primarily or exclusively by humans with minimal reliance upon machines. In essence, since privacy has to do with controlling the dissemination and use of Alice's personal information, techniques are distinguished at this first level according to how the control is actually implemented: are computers used to effect this control, or are human means used?

Levels 2 and 3. Within each of the above branches, technological and societal, privacy techniques are further classified according to whether Alice is doing something to effectively create new personal information, or a "holder" (an entity other than Alice who is in possession of Alice's personal information) is doing something with that information. More specifically, Alice may perform an action (on her local machine or on some remote machine) that is noticed by another entity Eve. Even if Eve does not know the content of the action (such as the data that was actually transmitted in a communications session between Alice and Bob), new personal information has been created if Eve becomes aware that this action is taking place. This personal information will typically include the participants, type, time, and location of the action in which Alice is engaged. As an alternative way for Alice to create new personal information, Alice may reveal specific attributes about herself (name and credit card information, for example) to another entity. From that entity's point of view, new personal information has been created if the entity did not previously know these attributes about Alice (*i.e.*, even though Alice has only disseminated some of her existing personal information, a new collection of personal information about Alice has been created at the other entity's site that did not exist previously).

Discussions about privacy often use the term "personal information" to refer only to attributes (data) about a subject, such as gender, address, salary, credit card number, political affiliation, health record, and the like. However, in many circumstances (such as corresponding with a specific individual/group, connecting to a certain Web site, or editing a particular file) the mere knowledge that the action was performed by Alice can also be personal information. This "action analysis" aspect of privacy is analogous to the "traffic analysis" aspect of security. This classification, therefore, recognizes both "actions" and "attributes" as valid types of personal information by categorizing privacy techniques according to which type of personal information Alice creates.

With respect to the "holder" (an entity in possession of Alice's personal information), there are also two possibilities. The holder may use this information for his or her own purposes, or may forward (disseminate) this information to some third entity. New personal information about Alice will be created in the latter case if the third entity did not previously know this forwarded information about Alice; but here it is the holder, rather than Alice, that is responsible for this creation.

Level 4. Privacy techniques can be further categorized according to the threat model under consideration. In particular, when Alice performs an operation, she may wish to hide this personal information from an intended target of the action (for example, she may wish to make an anonymous connection to a server), or she may wish to hide this personal information from unintended observers of her action (network eavesdroppers). When Alice reveals some personal attributes, again she may wish to hide it from the intended recipients of the data or from unintended recipients of the data. In the same way, a holder of Alice's personal information *I* may desire to protect against intended or unintended observers of his use of *I*, or against intended or unintended recipients when he disseminates *I*.

The classification up to this point is a balanced binary tree with sixteen leaves; see Figure 1 for the numerical labeling of these leaves (which will be referred to in subsequent text).

Level 5. (Not shown in Figure 1) We use the term "exposure" to refer to the unintentional release of information about the operations of Alice to other entities. Her activities are exposed (they are "brought to light" or "revealed") if an unintended entity can make the link between a particular action and the identity "Alice." That is, when Alice performs an operation, an *operation tuple* is formed: $\omega = (\iota, \alpha)$, where ι is the identity and α is the action. Private operations exist for Alice when unintended entities are unable to discover or infer the tuple ω.

Analogously, we use the term "disclosure" to refer to the unintentional release of the records of Alice to other entities. (A *record* may be narrowly defined in the sense of a single database record, or may be more broadly defined as a higher-level aggregation of information in a particular domain, such as a health record or a transaction record. In all cases, however, a record is some collection of information about a specific entity.) Alice's data is disclosed (given away or disseminated) if an unintended entity can make the link between one or more particular attributes and the identity "Alice." That is, when a record is created for Alice, a *record tuple* is formed: $\rho = (\iota, \overline{a})$, where ι is the identity and $\overline{a}$ is a collection of attributes. Private records exist for Alice when unintended entities are unable to discover or infer the tuple ρ.

Privacy techniques can be categorized at this level according to what kind of protection they offer for the two tuples ω and ρ. In particular, exposure-reducing transformations may be applied to operational data to increase the difficulty for an unintended party to construct ω, and disclosure-reducing transformations may be applied to record data to increase the difficulty for an unintended party to construct ρ. Protecting ω against exposure can be done in three ways: a transformation $\tau_\omega(\iota)$ may be applied to hide or remove the identity ι; a transformation $\tau_\omega(\alpha)$ may be applied to hide or remove the action α; or a transformation $\tau_\omega(\alpha)$ may be applied to hide the full tuple ω. Similarly, protecting ρ against disclosure can be done in three ways: a transformation $\tau_\rho(\iota)$ may be applied to hide or remove the identity ι; a transformation $\tau_\rho(\overline{a})$ may be applied to hide or remove the attributes $\overline{a}$; or a transformation $\tau_\rho(\rho)$ may be applied to hide the full tuple ρ.

5.1. Classification Summary

In more intuitive and descriptive terms, the classification proposed above is based on the discriminators that characterize any kind of investigation: who, what, when, where, why, and how.

- The "Why" of the classification is its title, its starting point: the classification's purpose is to categorize privacy techniques for online environments.
- "How" is the discriminator at Level 1: is privacy protected through technological means or through societal means?
- "Who" is the Level 2 discriminator: who is creating or using Alice's personal information *I* (is it Alice herself, or is it some other holder of *I*)?
- "What" is the Level 3 discriminator: what kind of personal information is being protected (is it the actions that Alice performs, or is it some attribute information about Alice; that is, does the privacy technique protect ω or does it protect ρ?
- "When" is the discriminator at Level 4: when is the information protected (as it is released to intended recipients, or as it is acquired by unintended recipients)?
- Finally, "Where" is the Level 5 discriminator: where is privacy protection applied (is it applied on the identity data, on the action/attribute data, or on the tuple)?

This set of discriminators makes the classification relatively simple to understand and to use.

*

6. EXAMPLES

USING THE NUMERICAL LABELS given in Figure 1 and the notion of privacy transformations from Level 5, the following examples illustrate the categorization of some well-known privacy techniques.

Label 1. There are circumstances in which Alice would like to have exposure privacy from an intended target of her operation. For example, she may wish to make anonymous or pseudonymous requests to a Web server and so will use techniques that provide privacy transformation $\tau_{\omega}(\iota)$ for her operation. Such techniques include Crowds[22] and anonymizers of various kinds,[23] as well as MIX networks[24] and onion routers[25] (if identifying information is omitted from the original embedded message). Alternatively, she may be comfortable with letting

22. See Michael K. Reiter & Aviel D. Rubin, "Crowds: Anonymity for Web Transactions" (1998) 1:1 ACM Transactions on Information and System Security 66, <http://avirubin.com/crowds.pdf>.
23. See *e.g.*, Anonymizer—Internet Privacy & Security Solutions, <http://www.anonymizer.com>.
24. David Chaum, "Untraceable Electronic Mail, Return Addresses, and Digital Pseudonyms" (1981) 24:2 Communications of the ACM 84, <http://gnunet.org/papers/p84-chaum.pdf>.
25. Michael Reed, Paul Syverson & David Goldschlag, "Anonymous Connections and Onion Routing" (1998) 16 IEEE Journal on Selected Areas in Communication, Special Issue on Copyright and Privacy Protection 482, <http://www.onion-router.net/Publications/JSAC-1998.pdf>.

the server know her identity, but may wish to hide the action she is performing and so will use techniques that provide privacy transformation $\tau_{\omega}(\alpha)$ for her operation. Such techniques include Private Information Retrieval (PIR),[26] in which, for example, Alice can request and receive records from a database server without the server knowing precisely which records she has requested. There do not seem to be practical situations in which Alice would desire to hide both her identity and her action from the target of the operation, but if such cases arise, Alice may use a combination of the above techniques.

Label 2. There are many circumstances in which Alice would like to have exposure privacy from unintended observers (eavesdroppers) of her operations. Privacy transformation $\tau_{\omega}(\iota)$ may be provided by various anonymizing services. Privacy transformation $\tau_{\omega}(\alpha)$ may be provided by techniques such as embedding Alice's actions in a continuous stream of faked transactions so that any observers are unable to construct an accurate profile of Alice or the group in which Alice is a member.[27] Finally, privacy transformation $\tau_{\omega}(\omega)$ may be provided by techniques such as onion routing and MIX networks. In all these transformations, if Alice wishes to reveal her identity, action, or both to the operation target (*i.e.*, if she wishes to hide this information *only* from unintended observers of her operation), then she would use encryption technology to protect an internal message containing the relevant information and would then operate on this encrypted message with an outer privacy-preserving transformation. Disabling third-party cookies on the Web browser is another simple example of a $\tau_{\omega}(\omega)$ technique.

Label 3. If Alice is deliberately revealing her personal attributes to another entity, there may be circumstances in which she would wish to do so without revealing her identity and so will use techniques that provide privacy transformation $\tau_{\rho}(\iota)$. For example, she may associate an anonym or pseudonym with this data, or she may ensure that the data is scrubbed of all identifying information prior to revealing it. Note that such a scrubbing operation must take into account the ways in which different attributes may interact to reveal identity.[28] If Alice wishes to hide the attributes themselves (without hiding her identity), privacy transformation $\tau_{\rho}(\vec{a})$ may be provided by mechanisms in privacy preserving data mining,[29] in which Alice can use randomized response

26. Dmitri Asonov, "Private Information Retrieval: An Overview and Current Trends" in *Proceedings of the ECDPvA Workshop, Informatik 2001, Vienna, Austria, September 2001,* <http://www.dbis.informatik.hu-berlin.de/fileadmin/research/papers/conferences/2001-gi_ocg-asonov.pdf>.

27. Yuval Elovici, Bracha Shapira & Adlai Maschiach, "A New Privacy Model for Hiding Group Interests While Accessing the Web" in *Proceedings of the ACM Workshop on Privacy in the Electronic Society, November 21, 2002,* pp. 63-70, <http://www.emeraldinsight.com/Insight/ViewContentServlet?Filename=Published/EmeraldFullTextArticle/Articles/1720150402.html>.

28. See *e.g.,* Latanya Sweeney, "Replacing Personally-Identifying Information in Medical Records: The Scrub System" in James Cimino, ed., *American Medical Informatics Association Proceedings,* Journal of the American Medical Informatics Association (Washington, DC: Hanley & Belfus, Inc, 1996) 333-337, <http://privacy.cs.cmu.edu/people/sweeney/scrubAMIA1.pdf>. See generally related privacy research at <http://privacy.cs.cmu.edu/people/sweeney>.

29, Wenliang Du & Zhijun Zhan, "Using Randomized Response Techniques for Privacy-Preserving Data Mining" in *Proceedings of The Ninth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining Held August 24-27, 2003,* <http://sai.syr.edu/facultypapers/Randomized%20Response%20Techniques.pdf>.

techniques to effectively hide her attributes from the intended recipient without compromising the accuracy of statistics computed over the entire population of users. Another example of $\tau_\rho(\vec{a})$ is oblivious polynomial evaluation,[30] in which Alice can give data to another entity Bob in such a way that Bob learns nothing about the data and yet is able to compute some function (a polynomial) on that data. As with Label 1, there do not seem to be practical situations in which Alice would desire to hide both her identity and her attributes from the intended receiver of her personal record ρ, but if such situations arise, a combination of the above techniques may be used.

Label 4. In all realistic situations, Alice would wish to hide her personal information from unintended recipients of this data, even as she reveals it to intended recipients. As with the case of performing an action, she may do this by encapsulating her identity, her attributes, or her tuple ρ in an outer layer that hides the desired information. Privacy transformations $\tau_\rho(\iota)$, $\tau_\rho(\vec{a})$, and $\tau_\rho(\rho)$ may therefore be provided by encryption technology. In particular, for $\tau_\rho(\iota)$ Alice may encrypt her identity using a key shared with the intended recipients and then anonymize or pseudonymize this encrypted blob using any appropriate source-hiding technique. For $\tau_\rho(\vec{a})$, in which Alice wishes to hide her attributes from unintended recipients but is not concerned with hiding her identity, secure channel technology such as Secure Sockets Layer (SSL)[31] is the most appropriate technique. The privacy transformation $\tau_\rho(\rho)$ is similar to $\tau_\rho(\iota)$, except that Alice encrypts both her identity and her attributes using the key she shares with the intended recipients before using a source-hiding technique on this encrypted blob. Anti-spyware and anti-adware software on Alice's machine also falls into the category of privacy techniques denoted by $\tau_\rho(\rho)$.

Label 5. Turning now from Alice to some other entity Harry that is a holder of Alice's personal information ρ, there are circumstances in which Harry would like to use ρ in such a way that there are intended observers of this use. For example, an organization may wish to change its business model or service offering in some way as a result of complaints from Alice or other data associated with Alice. However, Harry would like to hide Alice's identity, or her attributes, or both, from the intended observers. The transformations $\tau_\rho(\iota)$, $\tau_\rho(\vec{a})$, and $\tau_\rho(\rho)$ may be provided by inference control techniques. Such techniques are similar to ones used historically in a database context where the goal is to prevent inferences from being drawn from the association of separate pieces of stored data.[32] Here, the goal is to prevent observers from drawing inferences about Alice's personal information from the observed behaviour of Harry.

---

30. Yan-Cheng Chang & Chi-Jen Lu, "Oblivious Polynomial Evaluation and Oblivious Neural Learning" in C. Boyd, ed., *Advances in Cryptology—Proceedings of Asiacrypt 2001: 7th International Conference on the Theory and Application of Cryptology and Information Security Held December 9-13, 2001* (Springer-Verlag GmbH, 2001) 369-384, <http://www.iis.sinica.edu.tw/~cjlu/pub/ope.ps>.

31. See the OpenSSL Project <http://www.openssl.org>.

32. Dorothy Denning, *Cryptography and Data Security* (Reading, MA: Addison-Wesley, 1982). See especially chapter 6, "Inference Controls."

Label 6. There are circumstances in which the holder, Harry, of Alice's personal information ρ would like to use ρ without leaking anything about ρ to unintended observers of this use. Inference control techniques can again be used to provide the transformations $\tau_\rho(\iota)$, $\tau_\rho(\overline{a})$, and $\tau_\rho(\rho)$. However, in this case, any technology that helps to keep an organization's internal operations secret (*i.e.*, away from prying eyes) can also be a factor in maintaining the disclosure privacy of Alice's personal information. Thus, firewalls, anti-virus software, anti-spyware software, intrusion detection systems, and similar security products have a legitimate place within the classification of privacy techniques.

Label 7. Many situations exist in which Harry would like to disseminate Alice's personal information to intended recipients while hiding ι or $\overline{a}$ from these recipients. For example, a company may wish to reveal statistical information about its customer base (such as demographic information, buying patterns, or preference data) to other entities without revealing individual identities or attributes. Privacy transformation $\tau_\rho(\iota)$ may be provided by techniques for de-identifying data, as are used when the health records of a large number of patients are de-identified for release to organizations for statistical analysis or other research purposes; *k*-anonymity[33] is one example technique in this category. Privacy transformation $\tau_\rho(\overline{a})$ may be provided by data randomization techniques[34] in which attribute data is perturbed in such a way as to ensure with high probability that observations of individual attributes are incorrect but observations of population statistics are correct. Privacy transformation $\tau_\rho(\rho)$ does not seem to have a strong requirement in practice, but if there are uses for this, then a combination of de-identifying and data randomizing techniques may be applied.

Label 8. As with Label 4, in all realistic situations, Harry would wish to hide Alice's personal information ρ from unintended recipients of this data, even if he might need to reveal it to intended recipients. Again, privacy transformations $\tau_\rho(\iota)$, $\tau_\rho(\overline{a})$, and $\tau_\rho(\rho)$ may be provided by encryption technology (*i.e.*, ρ may be stored in an encrypted form on Harry's computer). In addition, as with Label 6, firewalls, anti-virus and anti-spyware software, intrusion detection systems, and similar products can be effective techniques for detecting unwanted intruders and preventing them from retrieving ρ from Harry's system. Another important class of techniques in this leaf of the classification is access control technology. If Harry can properly protect ρ through comprehensive access control tools and thoroughly-tested privacy policy enforcement architectures, then Alice's personal information will be effectively hidden from unintended recipients.

On the societal side of the classification, privacy techniques are again conceptually categorized according to what kind of protection they offer ω and ρ. In this branch of the tree, however, we find that the eight leaves (Label

33. Latanya Sweeney, "*k*-Anonymity: a Model for Protecting Privacy" (2002) 10:5 International Journal on Uncertainty, Fuzziness and Knowledge-based Systems 557, <http://privacy.cs.cmu.edu/people/sweeney/kanonymity.pdf>.

34. Huseyein Polat & Wenliang Du, "Privacy-Preserving Collaborative Filtering Using Randomized Pertubation Techniques" in *Proceedings of The Third IEEE International Conference on Data Mining Held November 12-22, 2003*, pp. 625-628, <http://www.cis.syr.edu/~wedu/Research/paper/icdm2003.pdf>.

9–Label 16) give rise, in practice, to only two classes of techniques, one of which is the trivial (NULL) class of no external party societal technique. In particular, for Labels 9 and 11, the only human means that Alice has to protect against exposure and disclosure is to be careful about engaging in activities and revealing her attributes. If she wishes to keep ω and ρ private from intended targets or recipients, she cannot generally rely on other human parties to assist in this process; the best she can do is to use discretion and to choose carefully regarding whether or not to perform any given operation or reveal any given attribute. This decision can be guided by some knowledge about the potential recipient entities, such as their organizational structure or procedures with respect to privacy, their accreditation or certification by recognized privacy agencies,[35] their adherence to government privacy guidelines, their compliance with privacy standards, their conformance with interoperability and "best practices" agreements for privacy, and the content of any privacy auditor's reports on their operations. In the final analysis, though, Alice has essentially only herself to rely upon to protect her exposure and disclosure privacy with respect to intended participants, whether she is engaged in some operation with an intended target or revealing some attribute data to an intended recipient.

Labels 10 and 12. For the remaining labels in this branch, the class of techniques available to Alice is defined by the existing legal infrastructure. In particular, for Labels 10 and 12 (in which Alice wishes to limit her exposure privacy with respect to unintended observers of her actions, and her disclosure privacy with respect to unintended recipients of attribute data that she reveals), government-initiated laws regarding the illegality of wiretapping assist Alice to protect her privacy. This protection may be in the form of a deterrent (a potential eavesdropper decides not to eavesdrop because of fear of the legal consequences of getting caught), or may be after-the-fact recourse or retribution when Alice's personal information has been compromised in this way.

Labels 13–16. The legal infrastructure for privacy protection is significantly more extensive for Labels 13–16 than it is for Labels 10 and 12. Legal mechanisms relevant to the holder of Alice's personal information may be classified into government-initiated law and contractual obligations. Government-initiated law governs or constrains a holder's handling of Alice's information (essentially independently of Alice herself) through the use of privacy-related laws, regulations, government-imposed guidelines, and so on.[36] Contractual obligations, on the other hand, arise from explicit or implicit contracts negotiated between Alice and the holder that define what the holder

35. See *e.g.*, Office of the Information and Privacy Commissioner of Ontario and Office of the Federal Privacy Commissioner of Australia, "Web Seals: A Review of Online Privacy Programs," *22nd International Conference on Privacy and Personal Data Protection Held September 2000,* <http://www.privacy.gov.au/publications/seals.html>.

36. See *e.g.*, OECD *Guidelines supra* note 13; *Children's Online Privacy Protection Act*, 15 U.S.C. 6501–6506, Pub. L. No. 105-277, Div. C, Title XIII, 112 Stat. 2681–2728 (1998), <http://uscode.house.gov/download/pls/15C91.txt>; American Institute of Certified Public Accountants, "Resources on International and US Federal & State Regulations," <http://infotech.aicpa.org/Resources/Privacy/>.

can do with Alice's information (for example, when Alice gives her address information to a store so that purchased items can subsequently be delivered); the legal infrastructure is brought into the picture only if a dispute or breach of contract occurs. As with Labels 10 and 12, there may be an element of deterrence here, but the primary protection offered to Alice is after-the-fact recourse and retribution.

*

7. IMPORTANCE AND USE OF A CLASSIFICATION

AS SUGGESTED IN THE INTRODUCTORY SECTION of this paper, a classification for privacy techniques is important because it allows different researchers and interested parties to have a common conceptual framework and common terminology in order to facilitate fruitful discussion and debate. Different technologies (both within a category and between categories) can more usefully be compared and contrasted to bring to light their strengths, weaknesses, and limitations with respect to protecting any particular aspect of Alice's privacy.

Equally importantly, however, a classification is useful because it can bring to our attention deficiencies in the suite of available technologies and thus suggest avenues for further research in the area of privacy protection. The classification proposed in this paper makes it clear in particular that there are areas on the societal side of the tree in which Alice has no real privacy protection other than her own judgment. Furthermore, even for the areas in which she has some legal protection, I note that in many cases of privacy violation there may be no adequate reparation for Alice. Once her personal information has been used or disseminated inappropriately by a holder, it may be that Alice cannot be sufficiently compensated for the damages suffered. The symmetry of the tree in this classification suggests the possibility that a deficiency in one area may be bolstered by a technique designed to achieve the same ultimate goal in the other half of the tree. Thus, for example, if Alice desires disclosure privacy with respect to the eavesdropping of her attributes by unintended recipients, she can rely not only on the deterrence and after-the-fact protection of relevant wiretap laws (Label 12), but also on a collection of technical mechanisms that fit within the corresponding technological branch (*i.e.*, Label 4). Note that within legal discourse a distinction is drawn between preventative approaches and remedial ones: tort provides remedy post-injury while other legal mechanisms, such as injunction, can be used preventatively. However, in practice, no legal approach is actually preventative. An injunction essentially says, "You are forbidden to do *X*, and if you do *X* these bad things will happen to you." But typically nothing physically prevents you from doing *X* if you are willing to accept the consequences, and so reliance must then be made on remedial approaches. Technical mechanisms, on the other hand, are more often preventative in practice: if you do not know the encryption key, for example, you are quite effectively prevented from seeing data that has been encrypted. Security practitioners have become familiar with the concept of "layers of security" for defense-in-depth in many environments; it may be wise for privacy practitioners

to similarly employ "layers of privacy" to provide strong privacy protection to individuals. The classification proposed in this paper can be a useful tool for determining which technologies are most complementary in this regard and can be used together in an effective way to achieve a particular privacy goal.

*

8. CONCLUSIONS

THIS PAPER HAS PROPOSED a classification for techniques that encourage, preserve, or enhance privacy in online environments. This classification is based on a set of discriminators that allows various techniques to be analysed and compared in meaningful ways. Furthermore, it provides a useful tool for determining which techniques can readily support each other to achieve a specific privacy goal.

It is hoped that this proposal will prove to be a useful foundation for further research in privacy. The existence of a comprehensive classification can lead to deeper understanding of the merits and deficiencies of particular privacy techniques, thereby stimulating further improvements in privacy protection for online users. I recognize, however, that an inherent shortcoming of any classification is the difficulty of obtaining a convincing proof of its completeness. How can one be assured that all possible instances will fit in a "natural" way within a proposed classification for a particular group of objects? With respect to the classification proposed in this paper, I have thus far found that the suggested categorization is relatively intuitive and, furthermore, I, and others, have been unable to find a privacy technique that does not fit comfortably into one of the defined categories. Nevertheless, continued examination and refinement of this classification is an important area of further research in which I encourage all interested parties to participate.