ANON: Towards a consolidated formal ontology for anonymity, unlinkability, and unobservability

Harry Halpin[1,*,†]

[1]Nym Technologies, 2 Place Numa-Droz, 2000 Neuchâtel, Switzerland

Abstract

As there had been widespread conceptual confusion in studies in terms of privacy and anonymity in privacy-enhancing technologies, Pfitzmann and Hansen's landmark "Anonymity, unlinkability, unobservability, pseudonymity, and identity management – A Consolidated Proposal for Terminology" provided the first comprehensive clarification. However, the definitions given were informal. Therefore, we present a sketch of the first formal ontology of this terminology using Semantic Web technologies.

Keywords

Anonymity, Unlinkability, Ontology, Semantic Web

1. Introduction

As opposed to previous work in formal ontology and knowledge engineering, privacy-enhancing technologies tend to focus on approaches backed by cryptography and information hiding using statistical techniques. Although the study of anonymity and cryptography was often conjoined in early work in the field [1], the fields diverged, and so unlike the mathematically precise security definitions used in provable cryptography in terms of security, the larger field of privacy has remained frustratingly vague and up for debate, with its proponents believing that privacy is not a precise and enforceable concept, and so (perhaps in bad faith) being willing to neglect privacy in the development of fields such as AI. One reason is there was never a commonly-accepted manner to formalize the privacy properties that privacy-enhancing technologies were hoping to achieve. Due to this, during the 1980s and 1990s there was widespread terminological confusion, from which emerged a commonly accepted terminology of privacy and anonymity: "Anonymity, unlinkability, unobservability, pseudonymity, and identity management – A Consolidated Proposal for Terminology" by Andreas Pfitzmann and Marit Hansen, originally put forward in 2000 by Pfitzmann and revised until 2008 with the help of Hansen.[1]

However, the work of Pfitzmann and Hansen has clear limits, as it is presented in an informal manner, with verbal appeals to set theory and probabilities [2].[2] This has made it difficult to use in practice when categorizing or reasoning about the properties of privacy-enhancing technologies [3]. Although only sketched, ANON is the first formal ontology that captures the many complex terminological distinctions made by Pfitzmann and Hansen [2] using Semantic Web technologies, in particular the Web Ontology Language (OWL). The ANON formal ontology is sketched in OWL2 and connected to each of the key terms of Pfitzmann and Hansen [2] in Section 3. We present next steps in developing ANON in Section 4, with the ultimate goal of encouraging knowledge engineering to take privacy and anonymity seriously as first-rate subjects of inquiry.

CEUR-WS.org/Vol-4340/paper12.pdf

The 25th International Conference on Knowledge Engineering and Knowledge Management (EKAW 2026) - Satellite Events (Workshops, Tutorials, Demos, Posters).

*Corresponding author.

harry@nymtech.net (H. Halpin)

0000-0003-2143-6965 (H. Halpin)

© 2026 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

[1]The cycle of revisions and refinements has sadly been halted by the death of Pfitzmann in 2010.

[2]An Appendix A1 was added to their terminology document that did present some simple class hierarchies for the terms [2].

CEUR Workshop Proceedings ceur-ws.org ISSN 1613-0073

published 2026-11-26

2. Background

Some of the earliest work on the intersection of formal ontology and privacy was having users define their privacy preferences using standards like the failed W3C P3P (Platform for Privacy Preferences) standard [4], which has shown itself amenable to being formalized [5]. This line of research evolved into the usage of ontologies to define access-control policies that could enforce privacy policies. However, what we are interested in demonstrating is not only the *use* of formal ontology to enforce some notion of privacy, but formally defining a semantic definition of privacy itself. While there exists a rich philosophical tradition, primarily from legal theory, of differing taxonomies of privacy [6], privacy tends to be difficult to formalize due to its dependence on social context, as exemplified by Nissenbaum's definition of privacy as contextual integrity [7]. Attempts have been made to formalize these fuzzy contextual notions of privacy [8], as well as laws like the General Data Protection Regulation that offer a clearly-defined structure for legal reasoning. Yet, privacy-enhancing technology has generally put forward *data minimization* – the inability to do data collection *technically* – as a solution rather than ensuring that the flows of data are "appropriate." Thus, the primary goal of privacy-enhancing technologies like Nym [9] and Tor [10] has been creating technical systems that can be proven not to collect data in the first place. This leads to the idea of data anonymization and so anonymity as one of the core tenets of privacy, where anonymity has often been thought of as a sort of extreme on the spectrum of privacy [11]. Being less context dependent, it seems like anonymity should be possible to formalize. For example, entropy-based methods can be used to measure anonymity [12]. So as a first step, we will have to move from the more difficult notion of privacy to the more limited – and so formalizable – concept of anonymity.

While there has recently been work trying to formalize anonymity using a provable security paradigm [13], this work appealed to an "indistinguishability" notion that incompletely maps to the terminology given by Pfitzmann [14]. Later work tries to extend the concepts of Pfitzmann using game-hopping proofs, in a similar fashion to proofs used in differential privacy [3]. While there was early work on providing a schema with semantics by the European Commission's Joint Research Centre, this work itself ended up confusing pseudonymity with anonymity and was limited by its use of a simple RDF schema [15]. Research on privacy in general was quite varied, with one strand investigating anonymity via modal logic [16] and another strand formalizing broader notions of identity [17]. Semantic Web ontologies such as DPV and PrOnto exist for data protection but not anonymity [18, 19]. However, none of these ontologies clearly focus on anonymity as used in privacy-enhancing technologies, whose canonical set of terms remains those defined by Pfitzmann and Hansen [2].

3. The ANON ontology

The ANON ontology formalizes the terminology of Pfitzmann and Hansen using the description logic *SROIQ*. Currently, the ANON OWL2 ontology is a work-in-progress and so lists the logical properties using description logic online at the following URL.[3] In this section, we will walk through each of the major classes of the ANON ontology and their relationships to each other. Note that a range restriction of role r to class B is represented as $r.A(B)$.

Communications technologies can be considered *systems* that provide various properties to their users, such as *anonymity*. The *surrounding* of the system is whatever is outside the system. The system and surrounding are the *universe*. First, there is an *attacker* who is observing *items of interest (IOI)*, a subset of all possible *items*, in a *system*. *Items* may be *messages* such as e-mails or files, as well as (possibly) human *subjects*, and so forth. An attacker may thus have one or more *observations* of an item of interest. Note for the concept of a *system*, there are items that are "inside" the system and "outside" the system, and so an attacker may be an *insider* that controls or can view parts of the system or an *outsider* that can view only interactions outside of a system. For any given *items of interest*, the attacker has one or more *observations*. The objective of observa-

[3] https://purl.org/anon

tions is to uniquely identify an item of interest if it is identical to a *subject* or *object*. A *subject* is a possibly acting *entity* such as, e.g. a *human being* (i.e. a natural person), a *legal person*, or a *computer*.[4]

$Universe \equiv System \sqcup Surrounding$
$System \equiv Universe \sqcap \neg Surrounding$
$Surrounding \equiv Universe \sqcap \neg System$
$Item \sqsubseteq Universe$
$IOI \sqsubseteq Item$
$IOI \equiv Entity \sqcup Action$
$Entity \equiv Actor \sqcup Actee \sqcup NaturalPerson \sqcup LegalPerson \sqcup Computer \sqcup Insider \sqcup Outsider$
$Attacker \equiv hasObservation.IOI \sqcap (\geq 1 hasAttacker.System)$
$Insider \equiv Attacker \sqcap System$
$Outsider \equiv Attacker \sqcap Surrounding$

Typically, there is a *sender* sending a message to a *recipient* using a *communication network* such as a VPN or an anonymity network such as the Tor onion routing network [10] or the Nym mixnet [9]. However, while the original ontology "assume[s] that each station is controlled by exactly one human being, its owner" but later versions assume the situation may be "more complex," so following Pfitzmann and Hansen we note that the sender and receiver are subjects and so may be *human beings* or they may be one or more *computers*. A message itself can then also be thought of as (non-exclusively) a *signal* or a *bitstring*. The goal of the attacker is to identify one or more items of interest with a particular *subject* based on their *observations*. The *perspective* of an attacker is the set of all observations that an attacker can carry out.

$Action \equiv \exists hasSending.Message \sqcup \exists hasReceiving.Entity$
$Sender \equiv \exists hasSending.Message \sqcap Actor$
$Receiver \equiv \exists hasReceiving.Entity \sqcap Actee$
$Message \equiv Signal \sqcup Bitstring$
$Subject \equiv Computer \sqcup NaturalPerson \sqcup LegalPerson$
$Perspective \equiv Attacker \sqcup \forall hasObservation.IOI$

Anonymity is when an attacker cannot identify[5] a particular subject with respect to more than one item of interest, given a set of observations. Note that this definition is also coherent with Shannon's classical definition of perfect secrecy. Thus, anonymity is identified as the property of a given item of interest only in relationship to an attacker.[6] Furthermore, each observation can be associated with a numerical value such as a probability (and observations may measure the knowledge gained by an attacker using a function over probabilities such as entropy), but these are outside the scope of description logic expressivity. However if we assume the description logic captures the state of a system only in a single moment in time, then we can attach a *probability* for that moment to each observation. With these caveats out of the way, we can then identify the **unlinkability** "of two or more items of interest (IOIs, e.g., subjects, messages, actions, ...) from an attacker's perspective means that within the system (comprising these and possibly other items), the attacker cannot sufficiently distinguish whether these IOIs are related or not"[2]. Being anonymous is therefore a class, but unlinkability is a property.

On this basis of defining unlinkability, we can then define an *anonymity set* as the set of possible items that cannot be uniquely linked to a given items of interest. This lets us define anonymity with regard to particular items of interest. As put by Pfitzmann and Hansen, "Anonymity of a subject means that the subject is not identifiable within a set of subjects, the anonymity set" [2]. For example "a sender s sends a message m (or set of messages) anonymously, iff s is anonymous within the set of potential

[4]A computer is called a "station" in earlier versions of Pfitzmann's terminology.
[5]As put by Pfitzmann and Hansen [2]: "not identifiable within" means "not uniquely characterized."
[6]As a limitation of our ontological approach, we do not deal with how robust anonymity is, i.e. how stable it is over time, as description logics cannot model time easily.

senders of *m*, the sender anonymity set of *m*." Anonymity can then be defined across senders, receivers, and relationships, so that one can have sender unlinkablility to a message, receiver unlinkability to a message, and relationship unlinkability between a sender and receiver. It should be noted that relationship unlinkability then means that each message is unlinkable to both a sender and a receiver, and so the unlinkability relationship is transitive, with regard to "this potentially communicating pair of subjects, each message is unlinkable" [2].

$IdentifiedIOI \equiv IOI \sqcap (= 1 anonymitySet.Attacker)$
$Anonymous \equiv \exists hasObservation.IOI \sqcap (\geq 2 anonymitySet.Attacker)$
$Probability \equiv \exists hasValue.Observation$
$Unlinkability \equiv \exists unlinkable.IOI \sqcap IOI$
$SenderUnlinkability \equiv \exists unlinkable.Message \sqcap Sender \sqcap Anonymous$
$ReceiverUnlinkability \equiv \exists unlinkable.Message \sqcap Receiver \sqcap Anonymous$
$RelationshipUnlinkability \equiv SenderUnlinkability \sqcap ReceiverUnlinkability$

Yet, there is another distinct notion that needs to be separated with, rather than confused with, anonymity: The notion of *undetectability*, where the "undetectability of an item of interest (IOI) from an attacker's perspective means that the attacker cannot sufficiently distinguish whether it exists or not" [2]. Although this may seem to be an impossible property to achieve, against certain attackers it may be possible. For example, a message sent disguised within an image may be undetectable to an attacker, even if they can observe the image, the attacker may not be able to observe the message. Of course, defining undetectability in a logical manner seems difficult, but it can be done by noting that for a particular item there are no observations by an attacker of the aforementioned item. Therefore, the relationship of undetectability is different from anonymity, because anonymity requires items of interest to be detectable and there to be a relationship between items of interest and one or more subjects.

$Undetectability \equiv IOI \sqcap \neg(\exists forObservation.IOI)$
$SenderUndetectability \equiv \exists undetectable.Message \sqcap Sender \sqcap Undetectability$
$ReceiverUndetectability \equiv \exists undetectable.Message \sqcap Receiver \sqcap Undetectability$
$RelationshipUndetectability \equiv SenderUndetectability \sqcap ReceiverUndetectabilty$

Since anonymity is defined in terms of observations and unlinkability given probabilities attached to those observations, note that "as the definition of undetectability stands, it has nothing to do with anonymity" insofar as the entire point of undetectability is that no observations are possible. Yet we can imagine undetectable systems that are also anonymous *in relation* to other components in the system, such as an undetectable message being sent anonymously. Therefore, Pfitzmann defines *unobservability* as the"unobservability of an item of interest (IOI)" which means that there exists "undetectability of the IOI against all subjects uninvolved in it and anonymity of the subject(s) involved in the IOI even against the other subject(s) involved in that IOI." This can be defined formally as follows:

$Unobservability \equiv Undetectability \sqcap Anonymous$
$SenderUnobservability \equiv \exists undetectable.Message \sqcap Sender \sqcap Unobservability$
$ReceiverUnobservability \equiv \exists undetectable.Message \sqcap Receiver \sqcap Unobservability$
$RelationshipUnobservability \equiv SenderUnobservability \sqcap ReceiverUnobservability$

Although this is only a small sketch of a much larger project, we have demonstrated that description logic can be used to formalize privacy properties. The ANON ontology is illustrated in Figure 3, although the figure is missing the less well-developed notion of unobservability. The ANON OWL 2.0 ontology contains 28 classes, 8 object properties, 17 defined (equivalent-class) axioms; HermiT classifies it as consistent with no unsatisfiable classes.

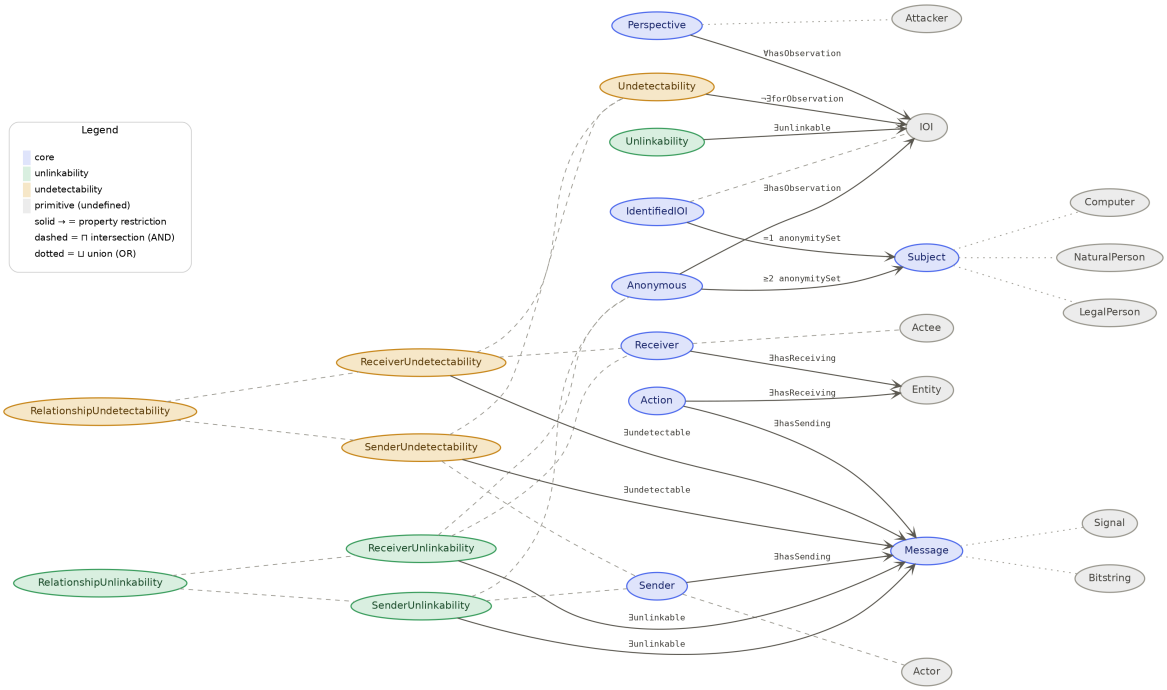


Figure 1: ANON Ontology Visualization

4. Conclusion

We have demonstrated that anonymity is a complex notion that can be broken down rigorously and even formally defined, at least in part. Our ANON ontology captures the definitions put forward by Pfitzmann and Hansen [2] for the core terms of anonymity and unlinkability. It also attempts, although more work is needed, to capture the more complex notions of unobservability and undetectability. Yet our treatment is not complete, insofar as it would have to be expanded to deal in more detail with unobservability and undetectability. Some relationships like pseudonymity are likely relationships between a particular individual and a particular name, and so an ABox relationship, and these relationships are not treated in ANON. Aspects of anonymity cannot be described easily in description logic, such as the concept of entropy and other functions over probabilities that can allow us to calculate whether a given attacker has been successful in identifying an item of interest with a given subject due to multiple observations over time. Identity management is also the least well-developed and least agreed-upon part of Pfitzmann and Hansen's terminological proposal, although prior work exists that could be integrated into ANON [15]. We cannot emphasize enough how this sketch is incomplete in terms of a larger ontology of privacy. There is a vast amount of hard work that can be done looking into the social (as well as legal) aspects of privacy, privacy as information flow, and privacy as part of AI safety and security. Focussing on anonymity first and formalizing it shows there are real notions here that should not be "hand-waved" away. Therefore, future work on ANON should focus on extending to more ABox-level relationships, defining identity management-related classes, relating ANON's relationship to real-world systems like Nym [9] and Tor [10], and the linking of anonymity to a larger comprehensive theory of privacy.

Declaration on Generative AI

The author(s) have not employed any Generative AI tools.

References

[1] D. L. Chaum, Untraceable electronic mail, return addresses, and digital pseudonyms, Communications of the ACM 24 (1981) 84–90.

[2] A. Pfitzmann, M. Hansen, A terminology for talking about privacy by data minimization: Anonymity, unlinkability, undetectability, unobservability, pseudonymity, and identity management, 2010. URL: http://dud.inf.tu-dresden.de/literatur/Anon_Terminology_v0.34.pdf.

[3] C. Kuhn, M. Beck, S. Schiffner, E. A. Jorswieck, T. Strufe, On privacy notions in anonymous communication., Proc. Priv. Enhancing Technol. 2019 (2019) 105–125.

[4] J. Reagle, L. F. Cranor, The platform for privacy preferences, Communications of the ACM 42 (1999) 48–55.

[5] T. Yu, N. Li, A. I. Antón, A formal semantics for P3P, in: Proceedings of the 2004 workshop on Secure web service, 2004, pp. 1–8.

[6] D. Solove, A taxonomy of privacy, University of Pennsylvania law review 154 (2006) 477–560.

[7] H. Nissenbaum, Privacy as contextual integrity, Wash. L. Rev. 79 (2004) 119.

[8] I. Kayes, A. Iamnitchi, Aegis: A semantic implementation of privacy as contextual integrity in social ecosystems, in: 2013 Eleventh Annual Conference on Privacy, Security and Trust, IEEE, 2013, pp. 88–97.

[9] C. Diaz, H. Halpin, A. Kiayias, The Nym network, 2021. URL: https://nym.com/nym-whitepaper.pdf.

[10] P. Syverson, R. Dingledine, N. Mathewson, Tor: The second-generation onion router, in: Usenix Security, 2004, pp. 303–320.

[11] S. Matthews, Anonymity and the social self, American Philosophical Quarterly 47 (2010) 351–363.

[12] A. Serjantov, G. Danezis, Towards an information theoretic metric for anonymity, in: International Workshop on Privacy Enhancing Technologies, Springer, 2002, pp. 41–53.

[13] M. Backes, A. Kate, P. Manoharan, S. Meiser, E. Mohammadi, AnoA: A framework for analyzing anonymous communication protocols, in: 2013 IEEE 26th Computer Security Foundations Symposium, IEEE, 2013, pp. 163–178.

[14] A. Hevia, D. Micciancio, An indistinguishability-based characterization of anonymous channels, in: International Symposium on Privacy Enhancing Technologies Symposium, Springer, 2008, pp. 24–43.

[15] G. Hogben, M. Wilikens, I. Vakalis, On the ontology of digital identification, in: OTM Confederated International Conferences: On the Move to Meaningful Internet Systems, Springer, 2003, pp. 579–593.

[16] Y. Tsukada, K. Mano, H. Sakurada, Y. Kawabe, Anonymity, privacy, onymity, and identity: A modal logic approach, in: 2009 International Conference on Computational Science and Engineering, volume 3, IEEE, 2009, pp. 42–51.

[17] M. Veeningen, B. De Weger, N. Zannone, Modeling identity-related properties and their privacy strength, in: International Workshop on Formal Aspects in Security and Trust, Springer, 2010, pp. 126–140.

[18] H. J. Pandit, B. Esteves, G. P. Krog, P. Ryan, D. Golpayegani, J. Flake, Data privacy vocabulary (DPV)–version 2.0, in: International Semantic Web Conference, Springer, 2024, pp. 171–193.

[19] M. Palmirani, M. Martoni, A. Rossi, C. Bartolini, L. Robaldo, Pronto: Privacy ontology for legal reasoning, in: International Conference on Electronic Government and the Information Systems Perspective, Springer, 2018, pp. 139–152.