



Three Dimensions of Privacy Policies

Victor Morel, Raúl Pardo

► To cite this version:

Victor Morel, Raúl Pardo. Three Dimensions of Privacy Policies. [Research Report] RR-9287, Inria - Research Centre Grenoble – Rhône-Alpes; CITI - CITI Centre of Innovation in Telecommunications and Integration of services. 2019. hal-02267641v3

HAL Id: hal-02267641

<https://inria.hal.science/hal-02267641v3>

Submitted on 22 Nov 2019 (v3), last revised 11 Sep 2020 (v4)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Three Dimensions of Privacy Policies

Victor Morel, Raúl Pardo

**RESEARCH
REPORT**

N° 9287

November 2019

Project-Teams Privatics



Three Dimensions of Privacy Policies

Victor Morel, Raúl Pardo

Project-Teams Privatics

Research Report n° 9287 — November 2019 — 43 pages

Abstract: *Privacy policies* are the main way to obtain information related to personal data collection and processing. Originally, privacy policies were presented as textual documents. However, the unsuitability of this format for the needs of today's society gave birth to others means of expression. In this survey, we systematically study the different means of expression of privacy policies. In doing so, we have identified three main categories, which we call *dimensions*, i.e., natural language, graphical and machine-readable privacy policies. Each of these dimensions focus on the particular needs of the communities they come from, i.e., law experts, organizations and privacy advocates, and academics, respectively. We then analyze the benefits and limitations of each dimension, and explain why solutions based on a single dimension do not cover the needs of other communities. Finally, we propose a new approach to expressing privacy policies which brings together the benefits of each dimension as an attempt to overcome their limitations.

Key-words: privacy policies, legal compliance, usability, enforcement

RESEARCH CENTRE
GRENOBLE – RHÔNE-ALPES

Inovallée
655 avenue de l'Europe Montbonnot
38334 Saint Ismier Cedex

Trois dimensions des politiques de protection de vie privée

Résumé : Les *politiques de protection de vie privée* sont le principal moyen d'obtenir de l'information liée à la collecte et au traitement de données à caractère personnel. Ces politiques étaient originellement présentées comme des documents textuels. Cependant, la non-convenance de ce format aux besoins de la société actuelle a donné lieu à d'autres moyens d'expression. Dans ce rapport, nous étudions de manière systématique les différents moyens d'expression des politiques de protection de vie privée. Ce faisant, nous identifions trois catégories principales que nous nommons *dimensions*, *i.e.*, les politiques en langage naturel, la représentation graphique des politiques, et les politiques lisibles par les machines. Chacune de ses dimensions se concentre sur les besoins spécifiques de la communauté dont elle est issue, *i.e.*, respectivement les juristes, les organisations et les défenseurs de la vie privée, et les universitaires. Nous analysons ensuite les avantages et les limites de chaque dimension, et nous expliquons en quoi les solutions basées sur une seule dimension ne couvrent pas les besoins des autres communautés. Enfin, nous proposons une nouvelle approche pour exprimer les politiques de protection de vie privée qui réunit les avantages de chaque dimension, dans le but de surmonter leurs limites.

Mots-clés : politiques de vie privée, conformité légale, utilisabilité, mise en application

1 Introduction

As of today, the main way to obtain information related to data collection and processing is through *privacy policies*. Privacy policies are typically presented as textual documents describing details such as data collection, processing, disclosure and management. Organizations collecting personal data (in what follows *data controllers*, or DC) commonly use privacy policies to inform individuals (in what follows *data subjects*, or DS) about how personal data is handled. DS are often required to read these policies — even though it rarely occurs [61] — or are at least presumed to do so and to decide whether they accept the conditions. Alternatively, giving DS the possibility of describing their own privacy policies has recently gained in popularity. This approach gives DS the time to reflect on their choices, and the possibility to consult experts and pairs. Nonetheless, privacy policies in their current format are hard to understand [22], for DS [73] as for experts [95]. In the sequel, we use *DS policies* to denote the privacy policies of individuals, and *DC policies* to refer to the privacy policies of organizations collecting personal data.

Requirements and recommendations to express privacy policies come from different sources such as privacy regulations, authorities and organizations. For instance, the General Data Protection Regulation (GDPR) [41] — the legal framework governing personal data collection and processing in Europe since May 2018 — requires more transparency for data processing from DC, and guidelines have been issued by the WP29¹ [122] to present their expectations. These requirements are necessary for privacy policies to be compliant with the legislation. Recommendations for drafting policies have also been made by different organizations to improve their readability. For example, the National Telecommunications and Information Administration [78] for mobile apps, and the WP29 [123] for IoT devices. Furthermore, authorities such as Data Protection Authorities (DPAs in the sequel) should be able to audit data processing systems, to ensure their compliance with the law and with the declared privacy policies.² All these requirements and recommendations can be summarized in three requirements:

- Privacy policies must be legally valid.
- Privacy policies must be understandable by all parties (including lay-users).
- Privacy policies must be effectively enforced through auditable mechanisms.

Existing methods to express privacy policies address some of these requirements, but not all of them. Different methods have arisen from different needs, and they target different audiences — from expert to lay-users. For instance, legal privacy policies are often written as long and complex documents which are necessary in court, but that are not easy to understand for lay-users. As an attempt to simplify these legal documents, organizations work on summarized versions of privacy policies or use visual aids to help users understand the risks of having their data collected. Another way forward in this direction is the use of techniques for extracting relevant information from existing privacy policies through Natural Language Processing (NLP) [120]. However, it is not powerful enough to be applicable to existing privacy policies yet. Ensuring that data is processed according to the requirements in privacy policies is not an easy task either. Some works — coming mostly from academia — propose an alternative format for privacy policies that can be read by computers. These works aim at bridging the gap between the textual legal requirements and their enforcement in the underlying system. Furthermore,

¹WP29 stands for Working Party 29, a European advisory board, now European Data Protection Board (EDPB)

²For instance, see the decision of the “Commission Nationale de l’Informatique et des Libertés” (CNIL), the French DPA, against Google LLC [26].

some of these proposals are equipped with auditing tools which facilitate, for DPAs or DS, verifying that no violations of a privacy policy have occurred. Unfortunately, these solutions are not widely used.

In this work, we analyze the state-of-the-art methods on expressing privacy policies. We provide a comprehensive picture of existing proposals in order to identify gaps and challenges. In doing so, we have identified that there exist three main ways to express privacy policies: natural language, graphical and machine-readable; which we call the *dimensions* of privacy policies. Each of these dimensions have arisen (independently) from different communities. Natural language comes from law experts, graphical from organizations and privacy advocates, and machine-readable from academics. Consequently, the content of this paper contextualizes knowledge often restricted to different communities that have been working in the same issue separately, and with different objectives. Unsurprisingly, each dimension mainly provides benefits to the specific community it was defined in. Therefore, we take the insights of our study to define how the different dimensions of privacy policies can complement each other. This synergy can include the benefits of each dimension and minimize their limitations. Hence, we propose a new type of privacy policies combining aspects from all dimensions, which we denote *multifaceted privacy policies*. More concretely, in bringing the above ideas to the forefront, our contributions are:

1. An in-depth study of the existing dimensions of privacy policies:
 - (a) Privacy policies expressed in natural language (Section 2). We denote these policies *natural language privacy policies*. Natural language privacy policies are necessary for legal compliance, and are required for DC to conduct lawful collection and processing of personal data.
 - (b) The expression of privacy policies graphically, *i.e.*, using visual aids such as icons and pictograms (Section 3). We denote these policies *graphical privacy policies*. Graphical privacy policies can be suitable for conveying intelligible information.
 - (c) Privacy policies that can be automatically processed by machines (Section 4). We denote this type of privacy policies *machine-readable privacy policies*. Machine-readability can be useful to provide tools to assist users and auditors in their tasks.
- For each dimension, we provide an overview of: i) its content; ii) the available tools; iii) its benefits; and iv) its limitations.
2. A categorization of existing works in each dimension according to a privacy taxonomy, and the specific features of each dimension.
 3. Insights from the study of existing means of expression of privacy policies (Section 5).
 - (a) Intrinsic limitations of mono-dimensional solutions.
 - (b) Guidelines to a multifaceted approach to express privacy policies, which overcomes the limitations of each dimension by combining their benefits.
 - (c) Missing elements of the privacy taxonomy in the privacy policies of different dimensions, thus shedding light on new paths for research in the design of privacy policy languages.

We discuss in Section 6 related work and conclusions.

2 Natural language privacy policies

Most legislations now require notices expressed in natural language to inform DS about the collection and processing of their personal data:³ the use of natural language is necessary to ensure that the policy has a legal value. The ways these documents can be authored — *i.e.*, drafted automatically or written manually — and the manners to assist their authoring can vary greatly. There are many ways to express privacy policies in natural language. In what follows, we present the content expressed by natural language privacy policies in Section 2.1, the tools used to assist their authoring and to analyze existing natural language privacy policies in Section 2.2, the benefits in Section 2.3, and the limitations in Section 2.4.

2.1 Content

Natural language privacy policies are familiar to the public as they have been adopted by a large range of online services such as social networks, file hosting services, or mobile applications *etc.* Because these privacy policies are expressed in natural language, they are not restricted in terms of content. This content can be categorized according to different taxonomies of privacy. In the following, we start by providing a high-level overview of existing taxonomies, and succinctly present our own.

2.1.1 Overview of existing taxonomies

In [106] Solove introduces one of the first, and perhaps the most known, taxonomy of privacy. This taxonomy focuses on activities that invade privacy, and distinguishes four categories: 1) *information collection* encompasses surveillance and interrogation, 2) *information processing* comprises aggregation, identification, insecurity, secondary use, and exclusion, 3) *information dissemination* covers breach of confidentiality, disclosure, exposure, increased accessibility, black-mail, appropriation and distortion, and 4) *invasions* includes intrusion and decisional interference. This taxonomy focuses on privacy harms, which makes it unsuitable for classifying the content of natural language privacy policies.

Paul *et al.* [84] introduced an evaluation framework to help DS assess how “privacy friendly” the privacy policies of IoT devices are. Concretely, they provide a scoring system based on the content of the privacy policy. The framework includes categories such as *Right to object*, *Right to access*, *Right to erase*, *Period of storage*, *etc.* As before, this work cannot be used to classify the content of natural language privacy policies, as it focuses on helping DS to take informed decisions based on the assessment of their privacy policies.

Wilson *et al.* [121] proposed a taxonomy tailored to privacy policies. It is composed of the following items:⁴ *First Party collection* “How and why a service provider collects user information”, *Third Party collection* “How user information may be shared with or collected by third parties”, *Access, Edit, Delete* “If and how users may access, edit, or delete their information”, *Data Retention* “How long user information is stored”, *Data Security* “How user information is protected”, *Specific Audiences* “Practices that pertain only to a specific group of users (e.g., children, Europeans, or California residents)”, *Do Not Track* “If and how Do Not Track signals for online tracking and advertising are honored”, *Policy Change* “If and how users will be informed about changes to the privacy policy”, *Other* “Additional sub-labels for introductory or general text, contact information, and practices not covered by the other categories”, and *Choice Control*

³For example, the GDPR mentions a list of information to provide where personal data is collected in Art. 13 and 14.

⁴We denote *item* a piece of information provided in a natural language privacy policy.

“Choices and control options available to users”. This taxonomy is appropriate for our purposes for two main reasons: 1) it was devised according to existing natural language privacy policies and therefore reflects their content; and 2) it encompasses most requirements of the current legislations or guidelines, such as the GDPR, the Fair Information Practice Principles (FIPPs) [44] in some cases, and the California Consumer Privacy Act of 2018 (CCPA) [107].⁵ This taxonomy was also used in Polisis [54] as discussed in Section 2.2.

2.1.2 Presentation of our taxonomy

We use a slight variation of Wilson *et al.*'s taxonomy, which does not change the content of the taxonomy but accommodates it to the purposes of our study. Concretely: i) we use *DS rights* to denote both *Access*, *Edit*, *Delete* and *Choice Control* as they relate to DS rights in the sense of the GDPR — see Chapter III of the GDPR: ii) subsume *Specific Audiences* and *Do Not Track* under *Other* as they are occasional items; and iii) we observe that a legal requirement is missing in the taxonomy, even though it is often found in natural language privacy policies: the *legal basis* of processing, we will therefore add it to our taxonomy. The differences are motivated by concern of a mapping with recent regulations (for DS rights and Legal basis) and practices (for Other).

Table 1 summarizes the chosen taxonomy, and shows what legal requirements appear explicitly in the GDPR, the FIPPs, or in the CCPA. We focus on the requirements of the GDPR, the FIPPs, and the CCPA as they are the three main texts (legislations or guidelines) that determine the content required when informing DS of data collection and processing. The GDPR is the text regulating personal data collection and processing in the EU, and many countries consider it since it has an extraterritorial scope. The FIPPs are guidelines designed by the United States Federal Trade Commission's (FTC) that represent widely-accepted principles concerning fair information practices. However, they have been considered outdated, notably due to the weight put on individual control instead of welfare [23]. The CCPA is a bill meant to enhance privacy rights and consumer protection for residents of California. It has been referred to “Almost GDPR in the US” [1]. We also consider other widely-known regulations, such as health data for HIPAA [116, Notice and Other Individual Rights], or children for COPPA [43, Â§ 312.4]). In what follows we examine the aforementioned items in detail, illustrated with examples from existing natural language privacy policies such as Facebook [42], Twitter [114], Dropbox [37], Netflix [79], and Google [51].

First Party Collection The most common item in natural language privacy policies is the first party collection, which describes *what* data is collected, *why* it is collected, and sometimes *how*. The type of data ranges from generic to more precise assertions, *e.g.*, respectively *we collect your data* and *your email address is collected*. Common types of data collected can be the name of the DS, an email address, geolocation, messages, *etc.*; or the social graph, more specifically to social networks. As an example, Facebook collects “Networks and connections. [...] information about the people, Pages, accounts, hashtags and groups you are connected to and how you interact with them across our Products[...]. We also collect contact information [...] (such as an address book [...]).” Cookies⁶ often have a distinct treatment, most likely because they are often collected by websites. It is common to find a dedicated paragraph for their management in a natural language privacy policy. Location data is often treated in a separate section as well because it can be collected from different sources — mobile applications, web browsers — or inferred from metadata — such as IP addresses. For instance, Twitter's privacy policy states: “Location Information: We require information about your signup and current location, which

⁵We discuss below the relevance of legislations with respect to the content of policies.

⁶Small pieces of data sent from a website and stored on the DS's computer by the DS's web browser [32].

	Description	GDPR	FIPPs	CCPA	HIPAA	COPPA
First Party collection	Type of data collected, purpose and collection mode.	●	●	◐	● _a	● _b
Third Party collection	Type of data collected, purpose and collection mode for third parties.	●	●	◐	◐ _a	● _b
Legal basis	Ground on which is determined the lawfulness of processing.	●	◐	○	○	○
DS rights	Rights of the DS, <i>e.g.</i> , right to access, to rectify, to port or erasure.	●	○	◐	● _a	● _b
Data Retention	Duration of data storage	●	○	○	○	◐ _b
Data Security	Modalities of protection of data, <i>e.g.</i> , encrypted communication and storage.	◐	●	○	● _a	◐ _b
Policy Change	Modalities of notification for policy changes.	◐	○	○	○	○
Other	Other items such as identity of DC, information related to DNT, to children ...	●/◐	●/◐	●/○	◐ _a	● _b

Table 1: Summary of our taxonomy, with the legal requirements of items. We use ● to denote *Required explicitly*; ◐ to denote *Addressed but not required*; and ○ to denote *Absent*. The subscript _a means that HIPAA only considers health data. The subscript _b means that COPPA only considers personal information from children, and notice must be addressed to parents.

we get from signals such as your IP address or device settings, to securely and reliably set up and maintain your account and to provide our services to you.” The purpose of processing often comes along the type of data. It is possible to find among the purposes, *marketing* and *advertising*, which are prevalent in natural language privacy policies. Analytics is often mentioned as a purpose to improve the functioning of services, to provide a better overview of what is actively used or not in a service, or to automatically retrieve malfunctions. Data can also be collected for security reasons: to remove illegal or harmful content, or to prevent payment fraud. Certain services collect data for research, and this broad purpose can be exempt of some constraints for the definition of a more concrete research purpose.⁷ DC can conduct data collection to operate a service: Facebook for instance mentions “Provide, personalize and improve our Products” as a purpose of processing, and this is a reason often put forward for data processing. It is also possible to find the collection mode in some natural language privacy policies: whether the data is collected automatically, by manual input of DS, or by any other mean.

Informing about the type of data, the purpose of processing, the recipients and the means of collection is required by the GDPR and the FIPPs. The CCPA gives the right to request first party collection, but does not make automatically mandatory. HIPAA requires to inform of “the ways in which the covered entity may use and disclose protected health information”. COPPA requires to inform of “what use, if any, the operator will make of the personal information collected”.

⁷See Recital 159 of the GDPR.

Third Party Collection Third Party Collection is a common item in natural language privacy policies, and it is therefore usual to find the third-parties to whom data will be transferred: they can be advertisers, or other business partners. The notion of *sharing* can also refer to other DS and subsidiary companies. It is usually composed of the same content as First Party Collection, *i.e.* type of data and purpose. For instance, Dropbox declares in its privacy policy: “Dropbox uses certain trusted third parties (such as providers of customer support and IT services) to help us provide, improve, protect and promote our Services. These third parties will only access your information to perform tasks on our behalf in compliance with this Privacy Policy, and we’ll remain responsible for their handling of your information per our instructions. For a list of trusted third parties that we use to process your personal information, please see our FAQ.”

Informing about third party collection is required by both the GDPR, the FIPPs, and COPPA.⁸ The CCPA gives the right to request categories of third parties, but does not require it except if data is sold to those third parties. HIPPA considers it implicitly.⁹

Legal basis Legal basis (or legal ground) is regularly found as a complement of the purpose of processing.¹⁰ A common legal basis for processing is consent, which consists, for DC, in retrieving an authorization from DS to legally collect their data. Consent has to be informed and specific under the GDPR,¹¹ and it still is often used as a legal basis. DC might consider the reading of their natural language privacy policies as a proper consent, without questioning the conditions to obtain consent [46]. Other legal basis can be found in natural language privacy policies, such as the necessity for the performance of a contract, compliance with legal obligations, protection of DS’s vital interests or public interest, and the legitimate interests of a DC. These legal basis are listed in the GDPR,¹² and major stakeholders generally consider cumulatively either all of them — such as Facebook which combines all possible legal bases — or a large subset — *e.g.*, Netflix’s policy considers all of them except public interest.

Informing about the legal basis is required by the GDPR, and not explicitly by the FIPPs which requires informing “whether the provision of the requested data is voluntary or required, and the consequences of a refusal to provide the requested information”. Legal basis information is not required by CCPA, HIPPA, and COPPA.

DS Rights DS can exercise rights regarding their data, and natural language privacy policies now often mention the rights to access, rectify, port and erase data, likely due to the influence of the GDPR.¹³ As an example, Google’s privacy policies mentions: “You can export a copy of your information or delete it from your Google Account at any time”. DS rights can be seen more restrictively as possibilities to opt-in or opt-out.¹⁴ Thus natural language privacy policies present how to subscribe or unsubscribe to specific services.

Informing about DS rights is required by the GDPR, HIPAA and COPPA, but not by the FIPPs. In CCPA, only the right to opt-out is explicit (other rights are ensured but Californians do not have to be explicitly informed of them in a privacy policy or in a privacy notice).

⁸Note that *recipient* in the sense of the GDPR encompasses first and third party collection.

⁹“The Privacy Rule covers a health care provider whether it electronically transmits these transactions directly or uses a billing service or other third party to do so on its behalf” [116, Who is Covered by the Privacy Rule].

¹⁰Article 13.1.c of the GDPR requires “The purposes of the processing for which the personal data are intended as well as the legal basis for the processing.” (Highlights from authors)

¹¹“Freely given, specific, informed and unambiguous indication of the data subject’s agreement” in Recital 32.

¹²See Art. 6.

¹³These rights are explicitly mentioned in the GDPR.

¹⁴Note that opt-out is now illegal in Europe.

Data Retention Natural language privacy policies often describe the period during which personal data will be stored. It can be a fixed value — *e.g.*, *30 days after data collection* — or a variable one — *e.g.* *as long as your account is active*. It often comes with the type of data, the purpose, and the legal basis of processing.

Informing about the retention time is required by the GDPR. COPPA addresses it but do not make it mandatory to inform about it. FIPPs, CCPA and HIPPA do not require it.

Data Security DC regularly explain in their policies how data is stored, if its communication is secured or its storage encrypted. As an example, Netflix’s privacy policy claims: “Security: We use reasonable administrative, logical, physical and managerial measures to safeguard your personal information against loss, theft and unauthorized access, use and modification. These measures are designed to provide a level of security appropriate to the risks of processing your personal information.”

Informing about the security of data is required by the FIPPs, but not by the GDPR although it mentions that “Personal data should be processed in a manner that ensures appropriate security and confidentiality of the personal data”.¹⁵ Similarly, COPPA states that “The operator must establish and maintain reasonable procedures to protect the confidentiality, security, and integrity of personal information collected from children.” HIPPA requires to inform of the “entity’s duties to protect privacy”. CCPA does not require to inform of this item.

Policy Changes The modalities of notification in the case of a change in the privacy policy can also be observed. Notification is usually by email or within the service’s interface, in some cases notifications are sent by regular mail or by phone.

This item is not required by the GDPR, the FIPPs, CCPA, HIPPA, nor by COPPA.

Other We subsume the identity and contact of DC, requirements towards specific audiences such children, and Do Not Track (DNT) under this item. The DC usually provides its identity, as well as its contact details if the DS has to lodge a complaint. In many legislations,¹⁶ children have specific considerations. As a result, it is possible to find a dedicated section in many natural language privacy policies, even if it is only to mention that personal data of children under thirteen is not collected without parental consent.

Identity of the DC is required by the GDPR, the FIPPs, and HIPAA, but not the rest of this item (although they specifically address data collected from children). CCPA more specifically require to “Make available to consumers two or more designated methods for submitting requests for information”, and do not address data collected from children. COPPA requires to inform of the identity and contact of DC, and is tailored to data collection related to children.

2.2 Tools

A variety of solutions exist to assist in the authoring of natural language privacy policies, ranging from the least to the most automated ones. We denote these tools *authoring tools*. We distinguish, in this section, *templates*, *generators*, and *retrievers*. While natural language does not limit the expressiveness of privacy policies, authoring tools are often tailored to websites and mobile app owners, and are constrained in terms of content. In addition to the authoring tools presented above, we denote *analysis tool* the piece of software able to parse or to analyze natural

¹⁵See Recital 39 of the GDPR.

¹⁶In particular Recitals 38 and 58 of the GDPR, and the Children’s Online Privacy Protection Act (COPPA).

language privacy policies, in order to produce a machine-readable or a graphical version of a policy. This section does not aim at providing an exhaustive list of the available tools.

Templates and generators Tools such as Docracy [36], Termsfeed [112], SEQ Legal [110], and 3DCart [3] provide a *fill-in-the-gap* form, where the author writes appropriate terms in the fields. We denote them *templates*. The redundancy is not taken into account, and no verification can be made regarding the validity of the terms written. It is not possible, for instance, to check if the email of a service owner is valid. More complex tools differ from *templates* by inputting information in a form using software components. We denote them *generators*. The author inputs information only once. Most *generators* do not allow incorrect data: email addresses without @ are highlighted, and the author do not have the possibility to go further in the process and generate the policy. *Generators* also give the option of expressing the same policy according to different legal contexts. A *generator* can propose text corresponding to the applicable legal framework. For instance, privacypolicies.com offers clauses specific to the GDPR or COPPA for an additional cost, if required. We can distinguish *light* from *detailed generators*. The former having a restricted set of parameters, the latter offering a choice between a policy tailored to websites or mobile applications, with a more exhaustive list of items [66]. Privacy Policy Generator [89], Privacy Policy Online [33], or GetTerms [49] are examples of the former, and PrivacyPolicies.com [91] and FreePrivacyPolicy.com [47] of the latter.

Retrievers Retrievers, such as those offered by Miao [75], Apolinarski *et al.* [10] and Yu *et al.* [125], automatically extract relevant information from code of mobile application, using static code analysis or user behavior analysis (for instance, in [10] the authors analyze sharing behavior when using online collaboration tools). These prototypes work on Android applications, in which personal data management are structured around the concept of permissions [50]. These permissions define the type of data accessible by an application, such as contacts, content of text messages, or Wi-Fi management. A *retriever* analyses those permissions, and interprets them according to well-defined rules to author a natural language privacy policy. In that case, an author does not necessarily have to input any information in addition to the code: the *retriever* can parse the name of the DC, the permission requested by a service or the third-party libraries, and can convert this information into natural language. However, *retrievers* are often tailored to a specific solution — mobile applications in most cases — and could thus be difficult to implement in other ecosystems. Furthermore, they cannot automatically retrieve certain information, such as the purpose of collection or the retention time. *Retrievers* reach the highest level of automation among authoring tools.

Analysis tools Analysis tools have been developed for over a decade. They focused on using Natural Language Processing or Information Extraction [27] to parse natural language privacy policies. An early work has been conducted by Brodie *et al.* [20]. The Usable Privacy Project lead by Sadeh [100] further investigated the automated classification of privacy policies. Within this project, Ammar *et al.* [8] conducted a pilot study for automatic text categorization. The Usable Privacy Project also developed a website privacy policy corpus [121], which will later be notably used by Polisis [54]. Zimmeck *et al.* [127] devise a hybrid solution combining machine learning classifiers (association rules) with crowdsourcing. Analysis tools achieve an accuracy averaging around 80%, which has not significantly improved since the first attempts.

2.3 Benefits

The main benefit of natural language privacy policies is their legal value. Most legislations require DC to provide a lawful statement detailing the processing of personal data, and natural language privacy policies often aim to fulfill this obligation.¹⁷

Legal value Natural language privacy policies are the only type of privacy policies with legal value as it is the standard format for legal texts.¹⁸ Lawyers rely on natural language to evaluate whether privacy policies are correctly drafted: these policies contain all details to determine whether there has been a violation. Also, lawyers use these policies to check compliance with data protection regulations, such as the GDPR. However, a document holding legal value is not necessarily compliant with the law. For instance, lawyers or DPAs may check that all items required by the legislation are provided to DS, and auditors can check that data processing is performed according to the policy. Legal compliance is twofold: with respect to information requirements, and with respect to the actual processing.

Value produced by authoring tools: The content produced by most authoring tools do not have legal value. Those tools do not provide legal advice, but rather general guidelines for policy authoring. These guidelines may be sufficient, but their legal validity is not guaranteed and should be verified by a lawyer. As an example, Iubenda advertises for its “Attorney-level compliance” [58], but advocates for a professional legal consultancy: “Nothing can substitute a professional legal consultancy in the drafting of your privacy policy” [59], and do not guarantee conformity with the law, which they claim “only a lawyer can do”. In other words: DC are responsible for the compliance with the law. They have to ensure that their privacy policies address all legal requirements, and to enforce the claims made in their policies.

2.4 Limitations

Ambiguity Natural language privacy policies can be ambiguous [94], as they may be interpreted in different ways. Reidenberg *et al.* [94, 96] presented privacy policies to privacy experts, law and policy researchers, who were ultimately unable to agree on some aspects of the policies. They proposed a crowd-sourcing annotation to tackle this issue, but admit that it would only provide a partial solution. This ambiguity is mainly due to the fact that a statement in natural languages can be interpreted in different ways. Ambiguity has a direct impact on the understanding, the enforcement, and the auditability of privacy policies. Ambiguity is also an explanation of the inaccuracy of analysis tools described in Section 2.2.

Understanding McDonald and Cranor [73] showed that it would take 200 hours a year for an average US citizen to read all the natural language privacy policies of the online services she used. This is clearly impractical, and thinking that DS read privacy policies before using a service is a *fiction juris*. All the more, nowadays, it seems highly inconvenient to spend a significant amount of time before using an online service.

¹⁷For instance, DC, in order to be compliant with the GDPR, must inform of the following: their identity and contact, the type of data collected, its purpose and its legal basis for processing, the recipient of data, the third-parties involved, the retention time, and the rights of the DS.

¹⁸It does not however mean that other dimensions do not have legal value, rather that natural language privacy policies are mandatory.

Enforcement & auditability Because they are currently ambiguous, natural language privacy policies are difficult to enforce: natural language lacks precise semantics, making it difficult to decide how data must be processed by the underlying system. Likewise, natural language privacy policies can also be hurdles to auditing: it can be difficult for an independent authority to compare stated and existing processing.

Summary

Natural language privacy policies are the most used medium to express privacy policies, and the tools used to assist their production can be categorized into *templates*, *generators*, and *retrievers*. These tools are often tailored to specific solutions, such as website or mobile applications, therefore restricting their scope. Analysis tools are not accurate enough to be trusted blindly. Natural language privacy policies are necessary for legal compliance, but suffer in practice from ambiguity and understandability.

3 Graphical privacy policies

In the previous section, we analyzed natural language privacy policies. They are necessary for legal compliance, but they can mislead DS when they attempt to read them, as they are often difficult to understand. As a consequence, other formats focused on DS understanding have been devised. Privacy policies can also be expressed with graphical representations, that we denote *graphical privacy policies*. Graphical privacy policies cover icons sets and standardized notices as well as solutions providing additional information, such as warnings or judgments, sometimes combined with simple text [98]. Graphical privacy policies often come from privacy advocates, but this is not only the case, notably since the WP29 explicitly mentioned icons as appropriate to convey privacy notices in their guidelines for transparency [122]. Privacy notices are means to inform DS, this term is heard in US context and is often understood as what we denote graphical privacy policies. We review in this section these means to express privacy policies graphically. In particular, we categorize each work based on: i) the elements in the taxonomy presented in Section 2 that it captures; ii) its features, whether it is made of icons, complementary text, or of something else; and iii) the intended audience of the language, *e.g.*, DS or DC. Table 2 in Section 3.4 summarizes our study.

3.1 Content

Based on their content, graphical privacy policies can be divided in three main types: *icons*, *standardized notices*, and *rating solutions*. Graphical privacy policies based on icons intend to express the content of privacy policies, for DC as for DS policies. Some of these icons try to cover all the items of the taxonomy introduced in Section 2. Other graphical privacy policies aim to express the same content as natural language privacy policies, but in a standardized and often comparable manner. Some graphical privacy policies provide rating information concerning certain aspects of privacy policies such as transparency level or potential risks. These solutions were not devised to meet the same requirements as natural language privacy policies, the content of these graphical privacy policies in that respect is often restricted. In the following, we describe the content of graphical privacy policies according to the elements of the taxonomy and their type (icons, standardized, or rating).

3.1.1 Sets of icons

The content of graphical privacy policies reviewed in this section lies in their icons, and sometimes in the simple explanations that comes with them. As an example, the symbol @ can represent collection of an email address, and a stylized calendar 📅 can represent retention time. But certain items are harder to express graphically. For instance, describing the legal basis of processing with the help of icons can easily be mistaken, and can mislead the intended audience instead of simplifying the understanding. In practice, graphical privacy policies have a restricted set of icons, and express specific items (seen in Section 2.1). Sets of icons initially emerged from academia, but they were quickly adopted by privacy advocates. A notable solution comes from the business sector.

Academia Sets of icons coming from academia do not sufficiently express the items presented in the taxonomy. Indeed, solutions such as Privicons [67] focus on informing mail correspondents of how the data should be handled instead (see Figure 1). Rundle [99] introduced the first set of icons in 2006 (see Figure 4). It includes icons for selling, and second-use of data, but the type of data cannot be specified. The PrimeLife project ([2, Chapter 15]) also proposed a set of privacy icons, that they tested in order to determine if they were understood. The icons presented hereafter (an excerpt is presented in Figure 2) can express retention time and sharing to third-parties, but the exhaustive set of icons is not available, making it difficult to judge the expressible content. Egelman *et al.* [38] developed a set of icons for the IoT, with the help of Intel first, then refined with crowdsourcing. The final set of icons (see Figure 3) focuses only on the type of data collected — voice, gesture, image — and its purposes — detection of gender, emotion, language. Recently, Rossi and Palmirani [97] proposed a Data Protection icon set, named DaPIS. The interesting features of their approach is that they based the icon set on an ontology named PrOnto, and they tested their set in order to refine it. Moreover, it stands out by emphasizing items recently introduced in the legislations, such as DS rights or legal bases for processing (see Figure 5). However, it does not consider the type of data.




Figure 1: Privicons



Figure 2: Excerpt of Primelife icons

Privacy advocacy Many privacy advocates contributed to this area and provided numerous sets of icons, such as Mehldau [74] who developed a set of 30 privacy icons, describing the type of data, third-parties, the purpose of processing and the retention time (see Figure 10). Generally speaking, icons within this subcategory can express more items, and items more related to natural language privacy policies. For instance, Aaron [4] came up with a set of icons that can only express three types of data, whether data may be disclosed to third-parties or not (see Figure 7) *etc.* Raskin [93] developed a set of icons for Mozilla. The type of data is not considered, and only the retention time, third-party use, ad networks, and law enforcement are considered (see Figure 9). Recently, a set of privacy icons was designed by Privacy Tech [90]. This set considers many types of data, as well as advanced representations of sharing, such as adequacy transfer (see Figure 6). It considers many common items, but not the rights of DS nor policy change — these items may be seen as less relevant for data transparency, even though mandatory under the GDPR.




Figure 3: “Is this thing on?” icons

	You agree not to use this data for marketing purposes.
	You agree not to trade or sell this data.
	You agree to submit to a third-party audit program on data use; if government has requested access to my data, you agree to involve my governmental ombudsman.
	You agree to make available to me the data that you have on me without my having to pay for it/at a minimal charge.
	You allow me to address inaccuracies in the data and request its removal.
	You agree to take reasonable steps to keep my data secure.
	You agree to arrange with X organization to help resolve any disputes we have over your treatment of this data. [The seal / name of the entity follows.]

Figure 4: Rundle set of privacy icons

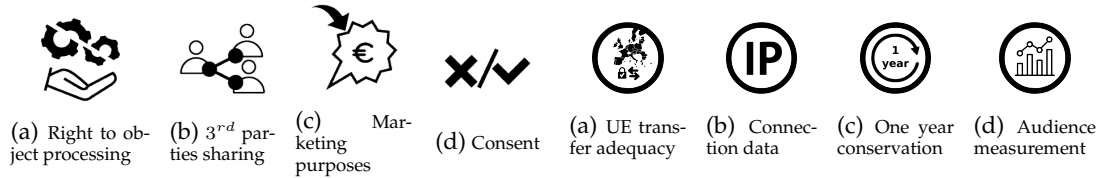


Figure 5: Excerpt of the DaPIS icon set

Figure 6: Excerpt of the Privacy Tech icons



Figure 7: Privacy Commons icons set

Business sector A notable example of privacy icons are the android permissions [50], created by Google. They present icons combined with simple natural language (see Figure 8). For each application installed on a mobile phone running Android, the permission manager presents a short graphical policy. Only little information is presented (the type of data collected, and processing in recent versions, but not the purpose for instance), and DS have to look into the natural language privacy policy in order to find more information.

3.1.2 Standardized notices

Another line of work considers the content described by the taxonomy, but as standardized notices instead of icons. These standardized notices are often represented in tables [64], but the key concept is the common vocabulary among notices. Kelley *et al.* [64] represent policies in a table such as nutrition labels observed on food packaging (see Figure 12). They developed a privacy nutrition label based on P3P, with the goal of providing efficient and well-organized privacy information. They present the fine-grained information in a table such as nutrition labels observed on food packaging. Polisis by Harkous *et al.* [54] can represent the natural language privacy policies as a combination of icons (see Figure 13a), highlights of the corresponding




Figure 8: Example of android permissions




Figure 9: Raskin's set for Mozilla




Figure 10: Mehldau's set of icons

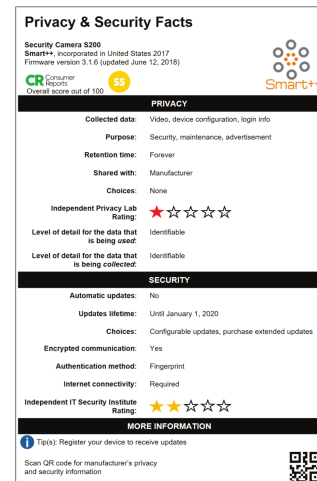


Figure 11: IoT label

paragraphs in the natural language privacy policies, and a flow diagram (see Figure 13b). Because Polisys relies on supervised machine-learning, *i.e.*, on a labeled corpus, it classifies natural language privacy policies in a standardized way. Emami-Naeini *et al.* [40] conduct a survey in order to rank the factors of IoT devices purchase. They determined that security and privacy were among the most important factors of purchase, and consequently developed an IoT privacy label to improve information visualization (see Figure 11). Cranor analyses the impact of the development of standardized mechanisms of notice and choice in [29], and more specifically the efforts conducted around P3P. Cranor reconsiders the advances made in standardization,¹⁹ as well the limitations, lack of adoption and enforcement.

3.1.3 Rating solutions

Certain graphical privacy policies do not consist of icons but provide other graphical representations instead. These solutions chose to present extra information related to privacy policies, often a judgment of the risk level associated to a DC policy, or a comparison between DS and DS

¹⁹Note that icons are considered as part of standardization efforts in [29].



Figure 12: Privacy Nutrition Label




Figure 13: Polisis

policies. We denote them *rating solutions*. Both academics and privacy advocates contributed to the solutions.

Academia Privacy Bird [30] is one of the first graphical tool. It consists of a colored bird, where the color indicates the matching (green for a match between the DS policy and the website's DC policy, red for conflict, yellow for uncertain, gray when disabled) (see Figure 14). It is represented as an add-on for Internet Explorer, restricted to Microsoft Windows. A dedicated website [30] provides an explanatory tour as well as a feature named *privacy finder*: Privacy Bird is then used as an indicator when browsing the web [21] [29]. Privacy finder displays the search results of a search engine, combined with the analysis of Privacy Bird. The bird was placed alongside search results, and was influential in the choice of shopping websites [113], notably when the items being purchased were likely to have privacy concerns [39]. DS could rank the results according to the matching between their DS policy and the websites DC policies. Van den Berg and Van der Hof [119] devised a wheel whose spokes show how data is handled (see Figure 15). Their solution highlights fairness of processing rather than transparency: it issues a

judgment on the processing, but shows little information with respect to what data is collected, by whom, and for what purpose. Sweeney *et al.* [109] proposes a simplified interface for access requirement to medical data called Datatags — further developed in [15] (see Figure 16). To each piece of data can be associated a tag presenting the risks, the security features associated, and the credentials required to access it. Hagan describes in [53] a user-centered privacy policy design project. The proposals include for instance a Visual Data Privacy Diagram to intuitively visualize data flow, Multi-character stories to present concrete situations (see Figure 17), and Context-Specific Alert for a selection of common questions regarding location data. The proposition cannot and do not intend to express items defined in Section 2, but attempts to increase DS awareness about consequences of data processing. In [82], Pardo & Le Métayer present a web interface to inform DS about the potential risks of their privacy policies. The interface is composed of a user-friendly form for DS to input their privacy policies and a set of risk analysis questions, *e.g.*, “Can company X collect my data?” (see Figure 18). DS simply need to click on “Analyze” to automatically obtain the answer to the questions. Additionally, DS may introduce risk assumptions in order to specify possible misbehaviors that the collecting parties can perform. In Section 4, we describe in detail the underlying privacy language and the automatic risk analysis.



Figure 14: Privacy Bird



Figure 15: Privacy wheel

Tag Type	Description	Security Features	Access Credentials
Blue	Public	Clear storage, Clear transmit	Open
Green	Controlled public	Clear storage, Clear transmit	Email- or OAuth Verified Registration
Yellow	Accountable	Clear storage, Encrypted transmit	Password, Registered, Approval, Click-through DUA
Orange	More accountable	Encrypted storage, Encrypted transmit	Password, Registered, Approval, Signed DUA
Red	Fully accountable	Encrypted storage, Encrypted transmit	Two-factor authentication, Approval, Signed DUA
Crimson	Maximally restricted	Multi-encrypted storage, Encrypted transmit	Two-factor authentication, Approval, Signed DUA

Figure 16: Final version of the Datatags

Privacy advocacy The ToS;DR initiative helps DS understanding the risks associated to a DC policy [111]. It started in 2011 during the Chaos Communication Camp. ToS;DR comprises not only icons, but also results from crowdsourced analyses in simple language. The idea of the project is to assess the data practices of web services by giving them badges, awarded by the project’s community. Once a service has enough badges to assess the level of protection of their terms for users, a class is automatically assigned by pondering the average scores (see Figure 19).




Figure 4: Privacy terms communicated through archetypal user stories, about privacy preferences, scenarios, and consequences.

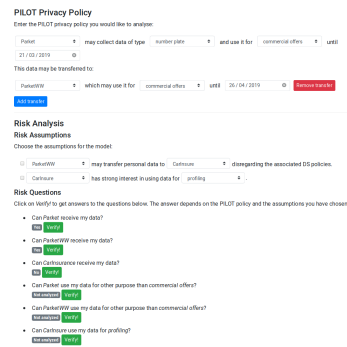


Figure 18: Input forms of risk analysis

Figure 17: Multi-character stories

- Class A** are the best terms of services: they treat you fairly, respect your rights and will not abuse your data.
- Class B** The terms of services are fair towards the user but they could be improved.
- Class C** The terms of service are okay but some issues need your consideration.
- Class D** The terms of service are very uneven or there are some important issues that need your attention.
- Class E** The terms of service raise very serious concerns.
- No Class Yet** We haven't sufficiently reviewed the terms yet.

Figure 19: Terms of Services; Didn't Read

3.2 Tools

Tools for representing graphical privacy policies (*graphical tools* in the sequel) are tailored to the web, and are often found as add-ons for web browsers.

Privacy Bird is represented as an add-on for Internet Explorer, restricted to Microsoft Windows. ToS;DR is also an add-on, for both Firefox and Chrome, as it ranks policies based on crowdsourced analyses by a community directly within the web browser. A website has been built to present Polisis [87], and add-ons for Chrome [85] and Firefox [86] are available (the add-ons redirect to the corresponding part of the website). The add-on "Disconnect Privacy Icons" [35], in collaboration with TRUSTe, which evolved from Raskin's set of icons for Mozilla (see Figure 9), provided an interactive and comprehensive view of privacy policies within the browser. The add-on would display icons according to a website privacy policy if the website complies with the solution.

All of these tools focus on the web, whereas the IoT is left unchallenged in that respect.

3.3 Benefits

Graphical privacy policies cannot be seen as legal commitment because they lack precise meaning, but they have other benefits: they can foster understanding.

Designed for lay-user understandability Many solutions coming from privacy advocates ([74], [93], and [4] for instance) aim to provide intelligible information to lay-users: "In order for privacy policies to have meaning for actual people" [92]. These solutions were built with

the will to popularize natural language privacy policies, and were designed to be understood quickly and take simplicity account. It is also the case for academic solutions such as the privacy labels [64], which “allows participants to find information more quickly and accurately”. Based on the principle that existing natural language privacy policies do not convey intelligible information about data collection and processing, Kelley *et al.* strove to provide a universal solution: “Our only requirement was that English be the participant’s native language”. Graphical privacy policies can also convey intelligible notices for scientists and physicians using sensitive datasets, such as the DataTags [109] — further developed in [15]. Their solution includes a simplified interface for access requirement to medical data, as this type of data is mostly restricted to medical practitioners and researchers.

Measuring understanding: Attempts were made to analyze what icons were recognizable and to measure their reliability. Egelman *et al.* [38] crowdsourced privacy indicators for the Internet of Things. In their study, they found out that some icons are well-recognized (for example, the camera symbol was recognized by more than 95% of participants as representing *video recording*), while others not (only 3.6% recognized the *voice command & control* icon). The PrimeLife project also proposed a set of privacy icons ([2, Chapter 15]), that they tested in order to determine if they were understood. They concluded that clear icons with few details were preferred. Kelley *et al.* [65] conducted a user study, to refine their privacy label. They compared the accuracy of information retrieval between their proposition and natural language privacy policies in natural language. As a result, they purposely combined simple natural language to prevent confusion, notably for the terms *opt-in* and *opt-out*.²⁰ Motti and Caine attempt in [77] to create a visual vocabulary for privacy by examining online images related to privacy. They analyze online images produced by users (such as Instagram), UI designers (*e.g.*, Material Design icons) and content producers (Shutterstock for instance). They identify five codes, *i.e.*, categories: action, objects, organizations, people, and abstract concepts. However, understanding of these icons and the vocabulary they represent is not tested. A promising attempt to measure understandability has been conducted by Rossi and Palmirani [97]. They performed three evaluations of their icon set in order to improve the recognition of icons. However, they regret the lack of diversity in the participants’ panel, notably for the educational level.

Legal Design: Legal design can be defined as “[the] application of design-thinking (processes by which design concepts are developed by designers) principles to the practice of law, to make legal systems, products, services and processes more useful, usable, understandable and engaging for all” [104]. For instance, the set of privacy icons from Privacy Tech [90] was designed with the GDPR in mind, to raise awareness among DS (see Figure 6). In a blog article, McCartney [72] discusses the Legal Design Lab lead by Margaret Hagan. She states that legal design — which includes graphical privacy policies *de facto* — could help people understanding complex legal issues. With respect to privacy, Hagan [53] proposed different tools to help lay-users understanding (see description in Section 3.1.3). Each proposition is an approach to convey privacy notices in a way that would provide better interaction with DS for a better understanding. Rossi *et al.* [98] proposed a set of Legal Design patterns to improve transparency in privacy notices. Drawing a parallel between consumer and privacy law, they devised of set of patterns, such as providing illustrating examples, or complementing privacy policies with “companion” icons.

²⁰Note that they also test the speed of retrieval, as well as comparisons between DC policies in addition to information retrieval.

3.4 Limitations

Ambiguity Though accessible to lay-users, graphical privacy policies may be interpreted in different ways, thus leading to ambiguities (see Section 3.1). The same icon can be interpreted in different ways according to the differences in culture, education level, or context *etc.* For instance, a euro symbol € can represent the commercial use of collected data, or that DS will be paid for having her data collected. Little has been done to produce a reasonably recognized set of icons for privacy — *e.g.* validated by a user study — despite the attempts of [38] and [2, Chapter 15] to see what were the most recognizable icons, and of [65] to provide a graphical policy where results could be found accurately. The three stages evaluation of Rossi and Palmirani [97] however leads the path to less ambiguous graphical privacy policies.

Incompleteness Graphical privacy policies are limited by their restricted scope. As seen in Section 3.1, existing graphical privacy policies are not as expressive as natural language privacy policies, due to the limited number of icons available. Some aspects are rarely mentioned, others only in complementary text and not in the graphical part of the policy. One aspects in particular is never mentioned in graphical privacy policies (policy change), and another is considered in one work only (legal basis).

Claim over legal compliance Some graphical privacy policies, such as cookie consent notices, have been used to claim legal compliance to retrieve consent. Degeling *et al.* [32] observed that a significant part (16%) of websites added cookie consent notices after the GDPR, but these notices do not always comply with transparency requirements according to Utz *et al.* [117] as they tend to use Dark Patterns to lure DS into giving their consent.²¹

Summary

Graphical privacy policies are promising for conveying summarized versions of natural language privacy policies, and they can rely on user-friendly tools to be adopted. However, they should come with explanations to ensure human understanding and mitigate their restricted content. See Table 2 for a visual and global overview of our study on graphical policies. Not all items of the taxonomy are considered: those not appearing in any of the surveyed works are omitted.

4 Machine-readable privacy policies

Many efforts have been devoted to the expression of *machine-readable privacy policies* — *i.e.*, privacy policies that can be automatically processed by computers. Most of these efforts were made by academics, and result in what has been called *privacy languages*. According to Kasem *et al.* [63], a privacy language is “a set of syntax and semantics that is used to express policies”. Many privacy languages have been proposed in the past twenty years (cf. [63, 118]). We review here the different ways in which privacy languages are used to express machine-readable privacy policies. In particular, we categorize each work based on: i) the elements in the taxonomy presented in Section 2 that it captures; ii) the type of enforcement mechanism it uses and whether it has been implemented; iii) additional tools for policy analysis or comparison; iv) the

²¹Dark Patterns are “instances where designers use their knowledge of human behavior (*e.g.*, psychology) and the desires of end users to implement deceptive functionality that is not in the user’s best interest” [52].

	Icons	Simple text	Rating	Standardized notice	DC	DS	1 st party	3 rd party	DS rights	Data retention	Data security	Legal basis
Privacy Bird [30]	✓ _b		✓		✓	✓	—	—	—	—	—	—
Rundle [99]	✓	✓				✓	●	○	●	○	●	○
Mehldau [74]	✓				✓		●	●	○	●	●	○
Privacy Commons [4]	✓				✓		●	●	○	○	○	○
Privacy Nutrition Label [64]		✓		✓	✓		●	●	○	○	○	○
Primelife [2, Chapter 15]	✓				✓		—	—	—	—	—	○
Raskin [93]	✓				✓		●	●	○	●	○	○
Privicons [67]	✓					✓	○	●	○	○	○	○
Privacy wheel [119]		✓	✓		✓		○	●	○	○	●	○
Android permissions [50]	✓	✓			✓		●	○	○	○	○	○
“Is this thing on?”[38]	✓				✓		●	○	○	○	○	○
Datatags [15]		✓	✓		✓		○	○	○	○	●	○
Hagan [53]		✓	✓		✓		●	●	●	●	○	○
Polisis [54]	✓	✓		✓	✓		● _a	● _a	● _a	● _a	● _a	○
Privacy Tech [90]	✓				✓		●	●	○	●	●	○
DaPIS [97]	✓				✓		●	●	●	○	○	●
IoT label [40]		✓		✓	✓		●	●	○	●	●	○
ToS DR [111]		✓	✓		✓		● _a	● _a	● _a	● _a	● _a	○
	Features				Type		Content					

Table 2: Categorization of graphical privacy policies

We use the subscript _a to denote that a solution extensively uses natural language in combination with graphical representations.

The subscript _b denotes that Privacy Bird uses icons, although not to specifically express items presented in the taxonomy, as opposed to other solutions within that category (see Section 3.1.1).

Features Whether the solution is made of icons or provides ratings about policies, and whether it provides explanation in simple natural language

Type (Type of policy) Whether the solution expresses a DC or a DS policy

Content Whether the solution can express the different items enumerated in Section 2. We use ● to denote that the solution can express most or all values; ● to denote that the solution expresses few values of the items, and is mostly insufficient; ○ to denote that the solution cannot express the item; and “—” to denote that the material does not permit judging whether the solution can express this item or not. Note some items of the taxonomy are omitted since no solution includes them.

intended audience of the language, *e.g.*, DS or DC; and v) whether it is intended to be directly used by lay-users. Table 3 in Section 4.4 summaries our study.

4.1 Content

We describe in this section the different types of content that machine-readable privacy policies include. This content is determined by the syntax of the privacy language. Many languages are defined using XML or JSON, and can therefore be automatically processed by machines. Other languages, however, are based on mathematical definitions (*e.g.*, logical languages), thus enabling the possibility of reasoning about them — these languages can easily be expressed in machine-readable formats due to the lack of ambiguity. Another important factor is the target audience of a language, *i.e.*, DC, DS or both. In what follows, we describe the content of machine-readable languages (according to the items defined in Section 2), the format used to express the policies and their target audience.

Access control languages such as XACML [9] and RBAC [101] have been among the first languages used for the specification of machine-readable privacy policies. Typically, these policies include the datatype to which they apply, and the set of entities with access privileges. Some extensions such as GeoXACML [70] include conditions depending on geolocation information, *e.g.*, “Alice can only access data from Lyon”. However, none of these languages captures concepts such as retention time, purpose or transfers that are in the privacy policy taxonomy described in Section 2. In other words, access control languages cannot impose any usage constraints after data has been accessed.

Usage control (UCON) [83, 88] appeared as an extension of access control to express how the data may be used after being accessed. To this end, it introduces *obligations*, which are actions to be executed after data has been received — *e.g.*, “do not transfer data item i to Company X” or “remove data on 28/01/2019”. These obligations make it possible to express items such as retention time, purpose and allowed data transfers. The Obligation Specification Language (OSL) [55] is an example of a fully-fledged UCON language together with an enforcement mechanism through Digital Right Management systems (DRMs) [45].

Neither access control nor UCON were developed with the idea of expressing privacy policies in mind. For instance, these languages do not offer mechanisms to describe DS policies. They are mostly used by DC to define their policies. New policy languages focused on expressing privacy policies appeared to address this problem.

Several languages dedicated to privacy policies have been proposed. A pioneer project in this area was the “Platform for Privacy Preferences” (P3P) [28]. P3P was conceived as a policy language for websites. It allows clients to declare their privacy preferences, and online service providers (mostly websites) to inform how they use customers’ data. P3P policies are specified in XML format, and include notions such as purpose, retention time and *conditions*. Conditions may be opt-in and/or opt-out choices for DS, or preferences based on enterprise data — *e.g.*, DS’s credit or service usage. Many extensions to P3P have been proposed [68, 12, 7], where its syntax has been extended — for instance, E-P3P [12] extends P3P’s syntax with obligations *à la* UCON. After P3P appeared, new languages with similar syntax have been proposed such as the “Enterprise Policy Authorization Language” (EPAL) [11], “An Accountability Policy Language” (A-PPL) [13], “Customer Profile Exchange” (CPExchange) [19], “Privacy Rights Markup Language” (PRML) [126], “Purpose-to-Use” (P2U) [60] and “Layered Privacy Language” (LPL) [48]. None of them add new features to the content of policies, but instead enhancements in terms of usability or enforcement (see Section 4.2). For instance, it is mandatory in LPL and PRML to include natural language explanations of policies, and EPAL offers automatic policy comparison.

Another line of work is that of formal privacy languages (*formal languages* in the sequel). S4P [17], SIMPL [69], QPDL [118], CI [80], PrivacyAPIs [71], PrivacyLFP [34] and PILOT [82] are languages which have their syntax and semantics defined mathematically. More precisely, they use formal languages such as *Linear Temporal Logic* [57], *First-Order Logic* [57] or *Authoriza-*

tion Logic [5]. However, not all of these formal languages have the same focus. S4P, SIMPL and PILOT are focused on expressing DS and DC policies. Thus, they do not differ much in content from the languages mentioned in the paragraphs above. It is possible to express types of data, conditions, purpose, retention time and allowed data transfers. Conditions are often more sophisticated than that of the languages mentioned above as they are based on logical languages. For instance, PILOT makes it possible to include spatio-temporal conditions which allow DS and DC to describe when, where, and by which devices data may be collected. On the other hand, CI, PrivacyAPIs and PrivacyLFP focus on encoding privacy regulations such as HIPAA [116], COPPA [43] or GLBA [115]. As a consequence, their expressive power is greater than languages focusing on DS and DC policies. They include temporal operators that make it possible to express policies about past and future events. For example, Barth *et al.* [16] express the following statement from COPPA “[...] an infant can only send identifiable information to a website, if her parent have previously sent their consent for data collection”. Finally, QPDL is a meta-language to reason about privacy languages. While privacy policies can be expressed in QPDL, it is not its intended use. The language was conceived as a framework to formally reason about different policy languages, *i.e.*, to compare the expressive power of different languages.

Jeeves [124] is a programming language with built-in support for a limited form of privacy policies. It allows programmers to declaratively specify confidentiality conditions based on the execution context. For instance, in a double-blind conference management system, paper authors can only be seen by organizers or the authors itself until the review process is completed.

4.2 Tools

In this section, we describe the mechanisms used to enforce machine-readable privacy policies, and existing tools to compare and perform analyses on policies for instance.

Formal Semantics Formal languages give meaning to their privacy policies by means of *formal semantics*. Typically, these semantics define what events may be executed depending on the privacy policies selected by the actors interacting in the system, *e.g.*, DS and DC. There are several ways to express semantics formally. For instance, SIMPL, S4P and CI use trace semantics, *i.e.*, they defined what are the allowed sequences of events (traces), given a set of privacy policies. PILOT uses small step operational semantics that define what events may be executed given the state of the system and the privacy policies of DS and DC. Jeeves, which is defined as a fully-fledged programming language with support for privacy policies, has its semantics formalized using lambda calculus [24]. Rei has its semantics defined as a set of logical rules in Prolog [25]. Though precise and unambiguous, formal semantics are not directly executable in most cases: there is a gap between a formal definition and the real implementation. Nevertheless, this gap may be very small, *e.g.*, Jeeves lambda calculus semantics were implemented as a Scala library, Rei’s semantics are encoded in Prolog, and PILOT semantics are implemented as a Promela model [56].

Informal Semantics Access control, UCON and privacy dedicated languages have their enforcement mechanisms specified as W3C specifications, specification languages such as UML, or they are simply implemented using a general purpose programming language. All these languages have in common that they use *request evaluation engines* to enforce privacy policies. Request evaluation engines take a *data request* and evaluate whether the requester may access the data based on the privacy policies. The content of data requests depends on the language. For instance, in RBAC, data requests contain the type of the requested data and the role of the requester. If the role of the requester matches one of the roles allowed by the policy associated

with the data, then data can be accessed. Usually, data requests include more information, *e.g.*, P3P data requests include data type, purpose of usage, requesting user, and the action to be performed (*e.g.*, read, write, delete, *etc.*). Most languages do not have mechanisms to enforce that data will be used according to the policies — *e.g.*, checking whether data is deleted before the retention time, or used for the specified purposes — but there are some exceptions. LPL erases automatically data from the central repository after the retention time has elapsed. UCON-based languages, such as OSL, use DRM to guarantee that obligations are enforced. A common factor of all these languages is that their request evaluation engines have been implemented and are ready to be deployed.

Policy comparison For some languages, algorithms have been devised to automatically compare policies. The goal is to determine, given two policies, which one is more restrictive. For example, a policy that allows data processing for research purposes during 7 days is more restrictive than a policy that allows data processing for advertisement and research during 90 days. Comparison is necessary to make it possible to mechanize consent. If a DC policy is more restrictive than a DS policy, then DS privacy preferences are satisfied. This step, although insufficient, is necessary for consent to be legally valid. Examples of such languages include EPAL, P3P and PILOT. In fact, the graphical tool Privacy Bird (see Section 3.2) uses P3P’s comparison algorithm to provide visual feedback to DS. CI, SIMPL and S4P follow a different approach. They define how restrictive a policy is, based on its semantics.²² Languages that do not distinguish DS and DC policies — such as RBAC, EPAL, A-PPL, or OSL — tend not to define algorithms to compare policies. This is not surprising, their goal is to enforce a policy typically defined by DC or system administrators.

Analysis tools Formal languages often come with tools to perform different types of automatic analyses. PILOT uses model-checking [14] to perform risk analysis. Given a DS policy and a set of risk assumptions, such as “Company X may transfer data to Company Y”, it is possible to automatically answer questions such as “Can Company Z use my data for advertisement?”. Rei comes with a Prolog interface where queries such as the above can be asked. PrivacyAPIs also uses model-checking to automatically verify properties about the privacy regulation HIPAA. It can for instance determine who can access patients medical files depending on their content or role.

4.3 Benefits

Machine-readable privacy policies have four main benefits: 1) they can be automatically enforced; 2) they can be audited; 3) it is possible to reason about their correctness; and 4) they make it possible to automate certain procedures. In what follows we explain each of these benefits in detail.

Enforcement As opposed to natural language or graphical policies, machine-readable policies can be automatically enforced. As described in Section 4.2, all policy languages have the means to guarantee that data is accessed according to the policies. Languages based on UCON or formal languages often provide stronger guarantees as they define how data are processed by all the parties after data collection. For example, they ensure that data is only used for purposes in the policies or that data is only transferred to allowed entities. Languages based on

²²Using trace semantics it is possible to compare policies based on the set of traces satisfying the policy. The less traces a policy satisfies, the more restrictive it is.

request evaluation offer weaker guarantees as they only protect access to the data, but not how the receiving party must process the data — only UCON offers limited support via DRMs. Nevertheless, due to their simplicity and ease of implementation, request evaluation languages are more widespread. Typically, every party holding personal data must implement the request evaluation engine. The implementation of formal languages tends to be more complicated. Normally, they require tracking actions applied on the data, or inferring what are the purposes for which data is used — as opposed to simply control access to data.

Auditability Machine-readable privacy policies enable the possibility of auditing whether data is being handled according to their respective privacy policies. This functionality is of great value for DPAs. Auditing mechanisms are typically implemented as logs that record the operations performed on sensitive data. For instance, EPAL requires to create an audit trail of access to keep track of whom has accessed personal data. In A-PPL, on the other hand, it is possible to specify *auditable operations* such as read or delete, and the enforcement records in a log every time that such operations occur. Ensuring the integrity of the logs is an orthogonal issue which is crucial for the legal validity of the auditing mechanism [18, 103].

Correctness The lack of ambiguity in policy languages makes it possible to precisely reason about their correctness, *i.e.*, that data is handled as stated in the privacy policies. This is specially true for formal languages. Their formal semantics can be used to formally prove certain correctness properties. For example, S4P, SIMPL and PILOT have been used to prove global properties such as “data is never used after its retention time”, or, “data is always used according to DS policies”. Moreover, languages focused on modeling privacy regulation — CI, PrivacyAPIs and PrivacyLFP — can be used to find inconsistencies in the regulation (if any). For example, it was possible using PrivacyAPIs to find unexpected ambiguities in HIPAA. These ambiguities were also found by commenters four years after it was enacted [71]. It is important to remark that there exists a gap between the formal semantics and its implementation — technical details not modeled in the semantics may lead to unforeseen violation of the properties. Therefore, formal languages should include auditing mechanisms, as the languages mentioned in the previous paragraph.

Automation Machine-readable privacy policies enable the possibility of automating certain procedures such as information communication and consent management. Automatic information communication facilitates transparency by making DS more aware of how their data is being handled — notably in ubiquitous systems where passive data collection is the norm. For instance, Das *et al.* [31] propose Personalized Privacy Assistants for the IoT. These assistants can inform DS of surrounding IoT devices thanks to the machine-readability of the information communicated. Automatic consent management can empower DS if managed in a protective way — *e.g.*, by mitigating the burden of choice [108] — and facilitate the retrieving of an informed consent for DC. Cunche *et al.* [76] devise a generic framework to manage informed consent in the IoT, using DS and DC policies based on PILOT semantics [82]. Automatic communication of privacy policies also makes possible a negotiation of privacy choices: DC and DS can interact more quickly by means of machines.

4.4 Limitations

The main limitations of machine-readable privacy policies are their lack of usability and adoption. As adoption relies among other things on human-understandability, understandable and

usable policies seems to be a condition *sine qua non* for their adoption.

Human understandability One of the most recurring criticism of machine-readable privacy policies is their lack of human understandability. Only a handful of languages such as XPref, SIMPL, LPL or PILOT take it into account: they include a natural language version of each policy. It is however questionable whether they can actually be understood. To put things into perspective, the OECD [81] conducted a study which shows that two third of adults from developed countries cannot conduct a medium-difficulty task related to ICT environments. Although privacy management was not mentioned in the OECD study, it is a medium-difficulty task, and solutions tackling privacy management must consider information-illiteracy. Machine-readable privacy policies should be expressed in languages close to natural language in order to be understood, or be complemented by friendly interfaces. Table 3 highlights the languages which address this issue in the column *usability*.

Lack of adoption Another pitfall for machine-readable privacy policies is their lack of adoption. It is arguably a consequence of poor human understandability. Most of the work done on privacy languages had few or no impact, apart from P3P. With the other solutions stemming from it (APPEL, E-P3P, ...) and the extension Privacy Bird for Internet Explorer, P3P obtained recognition out of the academic scope. It has been an official set of specifications of the W3C supported by the web browser Internet Explorer. Note that other languages were published as specifications by companies [19, 9] and can therefore be considered as having had some recognition. On the other hand, most formal languages lack a practical scalable implementation which makes it difficult to use in practice. Usability, implementation and widespread recognition are a rare combination in privacy languages.

Summary

Machine-readable privacy policies can provide means to express unambiguous privacy policies, and can be enforced as well as audited by authorities. However, they are often unintelligible for lay-users, which results in a lack of adoption. We provide a visual and global overview of machine-readable policies in Table 3. Not all items of the taxonomy are considered: those not appearing in any of the surveyed works are omitted.

5 Insights

In this section, we provide several insights that we identified as a result of our study. We show that each dimension is tailored to a specific audience, and that this is both 1) what makes it beneficial, but also 2) an obstacle to the compliance with all the requirements stated in Section 1 (*i.e.*, legal validity, understandability by all parties, and enforceability through auditable mechanisms). In Section 5.1, we highlight the benefits of each dimension for their particular audiences, and argue that a single dimension cannot comply with every requirement. In Section 5.2, we put in perspective the works which attempt to overcome the limitations of mono-dimensional solutions. Furthermore, we provide guidelines for a new approach to design privacy policies, which aims to cover the three dimensions. Finally, in Section 5.3, we discuss the coverage of the items in our taxonomy by privacy policies in the graphical and machine-readable dimensions.

	Usability	Syntax	Enforcement	Implemented	Tools	DS	DC	Time	Space	1 st party	3 rd party	DS rights	Data security	Data Retention
P3P [28]		XML	Informal	✓	Comparison	✓	✓	●	○	●	○	○	○	●
CPEXchange [19]		XML	Informal				✓	○	○	●	○	○	●	●
PRML [126]	✓	XML	Informal				✓	○	○	●	○	○	●	●
APPEL [68]	✓	XML	Informal	✓		✓	✓	●	○	●	○	○	○	●
E-P3P [12]		XML	Formal			✓	✓	●	○	●	○	○	●	●
Rei [62]		Formal	Formal		Analysis		✓	●	●	●	○	○	○	●
Xpref [7]	✓	XML	Informal	✓		✓		●	○	●	○	○	○	●
XACML [9]		XML	Informal	✓			✓	○	○	○	○	○	○	○
EPAL [11]		XML	Informal	✓	Comparison		✓	○	○	●	○	○	○	●
CI [16]		Formal	Formal			✓	✓	●	●	●	●	○	○	○
SIMPL [69]	✓	Formal	Formal			✓	✓	●	○	●	●	●	●	●
S4P [17]	✓	Formal	Formal			✓	✓	●	○	●	●	○	○	●
Jeeves [124]	✓	Formal	Formal	✓			✓	●	○	○	○	○	○	○
P2U [60]	✓	XML	Informal			✓		○	○	●	○	○	○	○
QPD [118]	✓	Formal	Formal			✓	✓	●	●	●	●	●	●	○
RBAC [101]		XML	Informal	✓			✓	○	○	○	○	○	○	○
OSL [55]		Formal	Formal	✓			✓	○	○	●	●	○	●	●
GeoXACML [70]		XML	Informal	✓			✓	○	○	○	○	○	○	○
A-PPL [13]		XML	Informal				✓	○	○	○	○	○	○	○
LPL [48]	✓	XML	Informal	✓		✓	✓	○	○	●	●	○	○	○
PrivacyAPIs [71]		Formal	Formal		Analysis	✓	✓	○	○	●	●	●	●	●
PrivacyLFP [34]		Formal	Formal			✓	✓	○	○	●	●	●	●	●
PILOT [82]	✓	Formal	Formal		Analysis	✓	✓	●	●	●	●	○	○	●
Features						Audience		Conditions		Content				

Table 3: Categorization of privacy languages.

Features Machine readable privacy policies specific features:

Usability Whether the language is *intended* to be understood by DS.

Syntax Whether the syntax of the language defined in XML or a formal language.

Enforcement Whether the language has a formally or informally defined enforcement.

Implemented Whether the language has been implemented.

Tools This column specifies type of available tools for the language.

Audience Whether the language can describe a DS or a DC policy.

Conditions Whether the languages supports conditional rules describing when (time) and/or where (space) data may be collected.

Content Whether the language can express the items described in Section 2.1. We use ● to denote that the item is explicitly included in the language; ○ to denote that the item is partially supported, *e.g.*, may be encoded through conditions or obligations; and ○ the item is not present in the language and cannot be encoded. Note some items of the taxonomy are omitted since no solution includes them.

5.1 Limitations of mono-dimensional solutions

Limitations in one dimension correspond to benefits in others. This is not a surprise, the different dimensions target different audiences and have different goals. Therefore, to each dimen-

sion corresponds benefits for an audience and limitations for another type of audience. In the following, we describe the limitations of each dimension for audiences outside the dimension.

Natural language privacy policies aim at defining the terms for data collection and processing. They must be precise enough so that, given a set of facts, a lawyer or a DPA can determine whether the privacy policy is consistent with what is properly enforced. In other words: they are required to check the compliance with the law and with the processing conducted by DC (see Section 2.3). Details specific to natural language privacy policies are often used by lawyers to check that the policy complies with privacy protection regulations — such as the GDPR — or they refer to functionalities of the system — *e.g.* logging or cookie management. In general, lay-users may not have the knowledge to fully understand these details, which makes it less accessible for them. For instance, the item *legal basis* privacy policies may be difficult to understand by DS. Likewise, natural language privacy policies do not include low level details related to the enforcement of the policies by a machine. In fact, those details are often unnecessary for the purposes of law enforcement and make it difficult for lawyers — who may lack the technical knowledge to understand the details — to use privacy policies. As a result, natural language privacy policies are mostly specified by DC, and they address DS and DPAs.

Graphical privacy policies aim at providing a simplified version of the policy to lay-users. They are useful for a better understanding. As a consequence, they are used by DC, and target DS. These policies do not contain details related to the legal aspects of the policy, nor aim to be automatically enforced by a machine. Graphical privacy policies are primarily used as complements to natural language privacy policies. Furthermore, their consistency with natural language privacy policies should be checked by DPAs in order to avoid misleading DS.

Machine-readable privacy policies, on the other hand, aim at being enforced by machines. They are written in a machine-readable format, and they include all the necessary details for the underlying system to enforce them. These details make them difficult to understand by humans, and are, consequently, unsuitable for lawyers and lay-users. Machine-readable privacy policies can be automatically enforced, and they enable robust auditing. Therefore, they are normally used by DC as part of their data collection and processing systems. Additionally, when they support DS and DC policies, they may be used for automatic policy comparison and consent management (for instance to palliate what Solove denotes as the structural problems of privacy self-management [105]).

Illustrative example In order to illustrate the differences in the details that the different dimensions capture, we use an example of privacy policies regarding retention time. We presented in Section 3 the icon  (from Privacy Tech Icons) that denotes that data is deleted after 1 year. We now show an excerpt of Facebook's privacy policy:

[...] when you search for something on Facebook, you can access and delete that query from within your search history at any time, but *the log of that search is deleted after six months.*

Facebook's policy is more precise than the icon: it refers to a very concrete piece of data which is produced after certain user action. However, these details may not be of prime interest for some lay-users, at least in a first stage. For instance, they may not know what a log entry or a query are. Hence this level of detail may be counter-productive for lay-users. Yet this information is required to determine whether Facebook is processing data according to the policy. Thus it cannot be omitted for legal purposes or for users who may be interested in more detailed information. Listing 1 shows the above policy in a machine readable language (APPEL-P3P). This policy includes details that are not present in the natural language policy above. For instance,

the format of the policy (XML) can be seen as including technical details about the underlying enforcement of the policy. Also, the policy includes additional technical parameters such as `xmins`, required so that the computer can retrieve the set of possible values for the element in the policy (the XML namespace), or the fact that retention time must be specified in days. Additionally, the enforcement mechanism of the policy would precisely define when the data is removed. For example, a background process checks on a daily basis whether the retention time has expired and, if so, deletes the data; or perhaps this process is executed once a year — both would comply with the privacy policies. The graphical and natural language versions of the policy only state for how long data will *not* be removed. These details are necessary for a machine to enforce the policy, but they have little or no value for lay-users and lawyers — in fact, some of these details negatively impact the understanding of the policy.

Intrinsic character of these limitations As a result, a single dimension cannot cover all the benefits required — *i.e.* legal validity, understandability by all parties, and enforceability. This is due to the tension between the details of a privacy policy and i) its suitability for lawyers, ii) lay-users, and iii) its automatic enforcement by machines. Concretely, there are details that only have meaning in one dimension and are irrelevant in others. Natural language privacy policies include details related to compliance with data protection regulations, which are unnecessary for the machine-readable dimension. Graphical privacy policies have the objective of being understood by a general audience, but this form of privacy policies have no use for lawyers or enforcement by machines.

```
...
<PURPOSE>
  <log-search-query />
</PURPOSE>
<RETENTION>
  <stated-purpose/>
  <EXTENSION>
    <retention-time days=182>
      <xmins="https://www.example.com/
        P3P/retention-time/" />
    </EXTENSION>
  </RETENTION>
...
```

Listing 1: Example of APPEL Policy.

5.2 Overcoming limitations

As argued in Section 5.1, a privacy policy expressed in only one dimension cannot satisfy requirements for lay-users, lawyers, and auditors. The current section aims to show that it is possible, however, to combine different dimensions to overcome their respective limitations. In many cases, limitations in one dimension can be addressed by other dimensions. For instance, natural language privacy policies may use graphical policies to enhance readability. Similarly, machine-readable privacy policies can use natural language privacy policies in order for the latter to be automatically enforced by DC. Furthermore, analysis tools provided with machine-readable privacy policies can be combined with graphical privacy policies to enhance the presentation of risks to DS. In other words: these three dimensions must be seen as complementary ways to express privacy policies, and they should be used together in order to meet the requirements stated in Section 1.

Several initiatives are already proposing cross-dimensions solutions, such as Harkous *et al.* [54] for natural language privacy policies and graphical privacy policies: they combined the accessibility of icons and simple text with the legal value of a natural language privacy policy. Policies can be more easily understood thanks to the results of the automatic analysis of natural language privacy policies. Le MÃ©tayer [69] proposed a combination of natural language privacy policies and machine-readable privacy policies: a machine-readable formal privacy language, enforceable, but close enough to natural language to be readable. PILOT [82] also combines machine-readable privacy policies and natural language privacy policies by providing a natural language user interface for DS to input their machine-readable privacy poli-

cies. Kelley *et al.* [64] added a graphical representation on top of P3P — a machine-readable privacy policy — resulting in both intelligible and enforceable privacy policies. Similarly, Rossi and Palmirani [97] based their icon set on an ontology, *i.e.*, they combined a graphical and a machine-readable approach. Icons are then machine-readable, even though they are not upheld by a formal semantics. These examples show that the combination of two dimensions makes it possible to take advantage of each dimension, without losing benefits.

Other initiatives attempted to dissociate the dimensions within the same solution. For instance, Van den Berg and Van der Hof [119] implemented a privacy wheel whose spokes present how data is handled. Those spokes have different layers: a first graphical layer aimed at DS, with the possibility to access two other layers of information comprising more details (“legal-istic” information). Another prominent example is the multi-layer approach of Abrams and Crompton [6]. They proposed a first layer as a short notice of whether data is collected, a second layer as a condensed notice of the data practices in a common graphic format, and finally a third layer, also called full notice, aimed at lawyers. However, these last two initiatives do not consider the machine-readability, and therefore the enforceability.

5.2.1 Recommendation: Multifaceted Privacy Policies

As we have seen, no existing solution encompasses the requirements for legal compliance, understandability, and enforceability. To mitigate this issue, we provide guidelines to define policies covering the three dimensions studied in this paper: natural language, graphical, and machine-readable. We denote these policies *multifaceted privacy policies*. We use the term *facet* to refer to a dimension of privacy policies when more than one is considered. The natural language facet can be verified for legal compliance, *i.e.*, lawyers should be able to check whether all information required by law is present. The graphical facet should be carefully designed to provide clear and concise information, and more information should be easily accessible. The machine-readable facet, endowed with a well-defined semantics, allows for the enforcement of the policy, and, in some cases, provides tools for analysis and comparison of privacy policies (see Section 4.2). Facets should be consistent: a graphical representation will surely omit details, but by no means it should mislead DS in their understanding of the privacy policy.

The natural language facet is mandatory, as it is a legal requirement. It can be decomposed according to the taxonomy in Section 2, and each element of the taxonomy can itself be further decomposed in simple *clauses*. A clause is a statement that would not make sense if further decomposed, *e.g.*, “we collect username”, or “collected data is stored until it is no longer necessary to provide our services”.

Figure 20 shows an example of multifaceted privacy policy for the clause highlighted in yellow. The clause refers to data transfers. Hence the multifaceted policy includes an icon from Raskin’s set of icons and a PILOT machine-readable policy as both can express data transfers. Note that the categorization of the solutions in each dimension according to the taxonomy presented in this paper can be of great use for the design of multifaceted privacy policies.

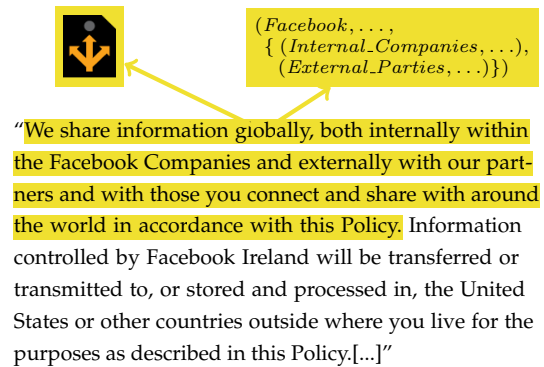


Figure 20: Multifaceted Privacy Policy Overview.

	Graphical Policies			Machine-Readable Policies			Legal Requirements		
	●	◐	○	●	◐	○	●	◐	○
1st party	33%	33%	34%	69%	31%	0%	80%	20%	0%
3rd party	17%	50%	33%	47%	26%	27%	60%	40%	0%
Legal basis	5%	0%	95%	0%	0%	100%	20%	20%	60%
DS Rights	22%	5%	73%	13%	17%	70%	60%	20%	20%
Data Retention	28%	11%	61%	30%	43%	27%	20%	20%	60%
Data Security	22%	22%	56%	17%	30%	53%	40%	40%	20%
Policy Change	0%	0%	100%	0%	0%	100%	0%	20%	80%

Figure 21: Coverage of taxonomy items by graphical and machine-readable privacy policies.

Ideally, multifaceted privacy policies should represent every clause in the graphical and in the machine-readable facets, in addition to the natural language facet, but it may not be necessary to include all facets for all clauses in practice. For example, the second clause in Figure 20: “Information controlled by Facebook Ireland will be transferred or transmitted to, or stored and processed in, the United States or other countries outside where you live for the purposes as described in this Policy.” corresponds to a specific detail which can be omitted in the graphical facet as it might confuse lay-users.

5.3 Missing taxonomy items

As we saw in the previous sections, not all items of our taxonomy are covered by existing solutions. In this section, we quantify the coverage of each item by the solutions we have surveyed. The aim of this section is to shed light on “forgotten” items, and hopefully guide future research on these directions.

Figure 21 summarizes our results. Each cell of the heat map shows the percentage of the surveyed tools (in a given dimension) that cover completely (●), partially (◐) or neither (○) an item of the taxonomy. For instance, the item 1st party is completely covered by 69% of policy languages (forth column, first row in Figure 21). Thus, it is the best covered item in the taxonomy. The right part of the table put the coverage in perspective of the legal requirements presented in Table 1.

The 1st and 3rd party taxonomy items have the best coverage in our study. Probably due to the fact that they express the most relevant information regarding data collection and processing for DS. These items are followed (in terms of coverage) by Data Retention and Data Security, which are absent from almost 50% of the languages we studied — except for Data Retention in machine-readable privacy policies. DS Rights is absent in around 70% of the graphical and machine readable policies languages. Most likely, because they refer to information difficult to express graphically, and they are outside of the realm of what machine-readable language are designed for. Nevertheless, work on expressing DS Rights graphically would be of great use for DS. Finally, legal basis and policy change are absent from all the works we surveyed. Possibly due to their absence in most legislations; 60% and 80% respectively.

6 Final discussion

Related work This work is not the first to propose an overview of the different manners to express privacy policies. Schaub *et al.* [102] present the requirements and best practices for presenting privacy “notices”. Their systematization of knowledge focuses on providing users impactful notices. In other words, they study how well they understand the messages conveyed by the privacy notices. Cranor [29] describes the notice and choice mechanisms, what P3P attempted to do to palliate issues raised by these principles, and why it failed in doing so. These works focus on the design of privacy policies to enhance usability for lay-users. In our work, we focus on connections of the graphical and machine-readable privacy policies with legal requirements (*i.e.*, natural language privacy policies). Although both these works consider the machine-readability of privacy notices, we highlight benefits that they did not consider, such as the possibility it offers for enforcement, auditing or automatic consent management.

To the best of our knowledge, this work is the first to propose a distinction between “dimensions” of privacy policies, and how these dimensions must be combined to provide legally valid, usable and enforceable privacy policies.

Conclusion In this paper, we have studied the different ways to express privacy policies: in natural language, with graphical representations, and using machine-readable means. We have categorized the existing works in each dimension according to a taxonomy of privacy policies, as well as their particular features. Additionally, we have studied the benefits and limitations of each dimension, and we have shown that the limitations of one dimension can be addressed by the benefits of the other dimensions. We have proposed a novel approach to express privacy policies combining the three dimensions, denoted multifaceted privacy policies, which overcomes the limitations of each dimension by bringing together their benefits. Finally, we have made explicit the degree of coverage of each element in the taxonomy by the surveyed works. Thus shedding light on future research directions for each dimension. We envision this work as an effort to bridge the gap between separate approaches — from the legal domain, design, and computer science — and to provide a big picture of how transparency can be ensured through privacy policies.

References

- [1] California's New Privacy Law: It's Almost GDPR in the US. URL <https://www.bankinfosecurity.com/californias-new-privacy-law-its-almost-gdpr-in-us-a-11149>.
- [2] Privacy and Identity Management for Life. URL <http://link.springer.com/10.1007/978-3-642-20317-6>.
- [3] 3DCart. Create an online store with 3dcart store builder. URL <https://www.3dcart.com/personalized-policy.html>.
- [4] Helton Aaron. Privacy Commons Icon Set .:aaron.helton:. URL <https://web-beta.archive.org/web/20090601215200/http://aaronhelton.wordpress.com/2009/02/20/privacy-commons-icon-set>.
- [5] Martín Abadi. Logic in access control. In *Proceedings of 18th IEEE Symposium on Logic in Computer Science (LICS 2003)*, 22-25 June 2003, Ottawa, Canada, page 228, 2003. doi: 10.1109/LICS.2003.1210062.
- [6] Marty Abrams and Malcolm Crompton. Multi-layered privacy notices: A better way. 2(1):1–4. URL <http://www.iispartners.com/downloads/Multi-LayeredNoticesPaperPublishedinPrivacyLawBulletinVol2No1June2005PLP2.1.orderform.pdf>.
- [7] Rakesh Agrawal, Jerry Kiernan, Ramakrishnan Srikant, and Yirong Xu. An XPath-based preference language for P3P. In *Proceedings of the 12th International Conference on World Wide Web*, pages 629–639. ACM. URL <http://dl.acm.org/citation.cfm?id=775241>.
- [8] Waleed Ammar, Shomir Wilson, Norman Sadeh, and Noah A. Smith. Automatic categorization of privacy policies: A pilot study. URL <http://repository.cmu.edu/lti/199/>.
- [9] Anne Anderson, Anthony Nadalin, B. Parducci, D. Engovatov, H. Lockhart, M. Kudo, P. Humenn, S. Godik, S. Anderson, S. Crocker, et al. Extensible access control markup language (xacml) version 1.0. URL <http://courses.cs.vt.edu/cs5204/fall105-kafura/Papers/Security/XACML-Specification.pdf>.
- [10] Wolfgang Apolinarski, Marcus Handte, and Pedro Jose Marron. Automating the Generation of Privacy Policies for Context-Sharing Applications. pages 73–80. IEEE. ISBN 978-1-4673-6654-0. doi: 10.1109/IE.2015.18. URL <http://ieeexplore.ieee.org/document/7194273/>.
- [11] Paul Ashley, Satoshi Hada, G  nter Karjoth, Calvin Powers, and Matthias Schunter. Enterprise privacy authorization language (EPAL). .
- [12] Paul Ashley, Satoshi Hada, G  nter Karjoth, and Matthias Schunter. E-P3P privacy policies and privacy authorization. In *Proceedings of the 2002 ACM Workshop on Privacy in the Electronic Society*, pages 103–109. ACM, . URL <http://dl.acm.org/citation.cfm?id=644538>.
- [13] Monir Azraoui, Kaoutar Elkhiaoui, Melek Onen, Karin Bernsmed, Anderson Santana de Oliveira, and Jakub Sendor. A-PPL: An Accountability Policy Language. In *Data*

- Privacy Management, Autonomous Spontaneous Security, and Security Assurance - 9th International Workshop, DPM 2014, 7th International Workshop, SETOP 2014, and 3rd International Workshop, QASA 2014, Wroclaw, Poland, September 10-11, 2014. Revised Selected Papers*, volume 8872 of *Lecture Notes in Computer Science*, pages 319–326.
- [14] Christel Baier and Joost-Pieter Katoen. *Principles of Model Checking*. MIT Press. ISBN 978-0-262-02649-9.
 - [15] Michael Bar-Sinai, Latanya Sweeney, and Merce Crosas. DataTags, data handling policy spaces and the tags language. In *Security and Privacy Workshops (SPW), 2016 IEEE*, pages 1–8. IEEE. URL <http://ieeexplore.ieee.org/abstract/document/7527746/>.
 - [16] Adam Barth, Anupam Datta, John C. Mitchell, and Helen Nissenbaum. Privacy and contextual integrity: Framework and applications. In *2006 IEEE Symposium on Security and Privacy (S&P'06)*, pages 15–pp. IEEE. URL http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=1624011.
 - [17] Moritz Y. Becker, Alexander Malkis, and Laurent Bussard. S4P: A generic language for specifying privacy preferences and policies. URL <http://www.msr-waypoint.com/pubs/122108/main.pdf>.
 - [18] Giampaolo Bella, Rosario Giustolisi, and Salvatore Riccobene. Enforcing privacy in e-commerce by balancing anonymity and trust. 30(8):705–718. ISSN 01674048. doi: 10.1016/j.cose.2011.08.005. URL <http://linkinghub.elsevier.com/retrieve/pii/S0167404811001052>.
 - [19] Kathy Bohrer and Bobby Holland. *Customer Profile Exchange (Cpexchange) Specification*. URL http://mail.ctiforum.com/standard/standard/www.cpexchange.org/cpexchangev1_0F.pdf.
 - [20] Carolyn A. Brodie, Clare-Marie Karat, and John Karat. An empirical study of natural language parsing of privacy policy rules using the SPARCLE policy workbench. In *Proceedings of the Second Symposium on Usable Privacy and Security - SOUPS '06*, page 8. ACM Press. ISBN 978-1-59593-448-2. doi: 10/b3tswp. URL <http://portal.acm.org/citation.cfm?doid=1143120.1143123>.
 - [21] Simon Byers, Lorrie Faith Cranor, Dave Kormann, and Patrick McDaniel. Searching for Privacy: Design and Implementation of a P3P-Enabled Search Engine. In David Martin and Andrei Serjantov, editors, *Privacy Enhancing Technologies*, volume 3424, pages 314–328. Springer Berlin Heidelberg. ISBN 978-3-540-26203-9 978-3-540-31960-3. doi: 10.1007/11423409_20. URL http://link.springer.com/10.1007/11423409_20.
 - [22] F. H. Cate. The Limits of Notice and Choice. 8(2):59–62, . ISSN 1540-7993. doi: 10/cgjkcd.
 - [23] Fred H Cate. The Failure of Fair Information Practice Principles. page 38, .
 - [24] Alonzo Church. An unsolvable problem of elementary number theory. *American Journal of Mathematics*, 58(2):345–363, 1936. ISSN 00029327, 10806377.
 - [25] William F. Clocksin and Christopher S. Mellish. *Programming in Prolog (4. ed.)*. Springer, 1994. ISBN 978-3-540-58350-9.

- [26] CNIL. The CNIL's restricted committee imposes a financial penalty of 50 Million euros against GOOGLE LLC. URL <https://www.cnil.fr/en/cnils-restricted-committee-imposes-financial-penalty-50-million-euros-against-google-llc>.
- [27] Elisa Costante, Yuanhao Sun, Milan Petković, and Jerry den Hartog. A machine learning solution to assess privacy policy completeness:(short paper). In *Proceedings of the 2012 ACM Workshop on Privacy in the Electronic Society*, pages 91–96. ACM. URL <http://dl.acm.org/citation.cfm?id=2381979>.
- [28] Lorrie Cranor, Marc Langheinrich, Massimo Marchiori, Martin Presler-Marshall, and Joseph Reagle. The platform for privacy preferences 1.0 (P3P1.0) specification. 16. URL <https://elearn.inf.tu-dresden.de/hades/teleseminare/wise0405/Act.%208%20Models%20Languages%20Pierangela/Materials/P3P.pdf>.
- [29] Lorrie Faith Cranor. Necessary but not sufficient: Standardized mechanisms for privacy notice and choice. 10:273. URL http://heinonline.org/hol-cgi-bin/get_pdf.cgi?handle=hein.journals/jtelhtell10§ion=22.
- [30] CyLab Usable Privacy and Security Laboratory. Privacy Bird. URL <http://www.privacybird.org/>.
- [31] Anupam Das, Martin Degeling, Daniel Smullen, and Norman Sadeh. Personalized Privacy Assistants for the Internet of Things. 2018. doi: 10.1109/MPRV.2018.03367733.
- [32] Martin Degeling, Christine Utz, Christopher Lentzsch, Henry Hosseini, Florian Schaub, and Thorsten Holz. We Value Your Privacy ... Now Take Some Cookies: Measuring the GDPR's Impact on Web Privacy. doi: 10/gfxgxm. URL <http://arxiv.org/abs/1808.05096>.
- [33] Daniel DelPercio. Privacy Policy Online (2011). URL <http://www.PrivacyPolicyOnline.com>.
- [34] Henry DeYoung, Deepak Garg, Limin Jia, Dilsun Kirli Kaynar, and Anupam Datta. Experiences in the Logical Specification of the HIPAA and GLBA Privacy Laws. In *Proceedings of the 2010 ACM Workshop on Privacy in the Electronic Society, WPES 2010, Chicago, Illinois, USA, October 4, 2010*, pages 73–82.
- [35] Disconnect. Privacy Icons. URL <https://web.archive.org/web/20160304013156/https://disconnect.me/icons>.
- [36] Docracy. An open source privacy policy for mobile apps. URL <https://web.archive.org/web/20171124185357/https://blog.docracy.com/post/27931026976/an-open-source-privacy-policy-for-mobile-apps>.
- [37] Dropbox. Dropbox Privacy Policy. Effective: 25 May 2018. URL <https://www.dropbox.com/privacy>.
- [38] Serge Egelman, Raghudeep Kannavara, and Richard Chow. Is This Thing On?: Crowdsourcing Privacy Indicators for Ubiquitous Sensing Platforms. pages 1669–1678. ACM Press, . ISBN 978-1-4503-3145-6. doi: 10.1145/2702123.2702251. URL <http://dl.acm.org/citation.cfm?doid=2702123.2702251>.

- [39] Serge Egelman, Janice Tsai, Lorrie Faith Cranor, and Alessandro Acquisti. Timing is everything?: The effects of timing and placement of online privacy indicators. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 319–328. ACM, . URL <http://dl.acm.org/citation.cfm?id=1518752>.
- [40] Pardis Emami-Naeini, Henry Dixon, Yuvraj Agarwal, and Lorrie Faith Cranor. Exploring How Privacy and Security Factor into IoT Device Purchase Behavior. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems - CHI '19*, pages 1–12. ACM Press. ISBN 978-1-4503-5970-2. doi: 10/gf5d6v. URL <http://dl.acm.org/citation.cfm?doid=3290605.3300764>.
- [41] European Parliament. General Data Protection Regulation.
- [42] Facebook. Facebook Data Policy. Date of last revision: April 19, 2018. URL <https://www.facebook.com/privacy/explanation>.
- [43] Federal Trade Commission. Children’s Online Privacy Protection Rule; Final Rule, . URL <https://www.ftc.gov/system/files/documents/federal-register-notice/2013/01/2012-31341.pdf>.
- [44] Federal Trade Commission. FAIR INFORMATION PRACTICES IN THE ELECTRONIC MARKETPLACE. page 208, .
- [45] Joan Feigenbaum, Michael J. Freedman, Tomas Sander, and Adam Shostack. Privacy engineering for digital rights management systems. In *Security and Privacy in Digital Rights Management, ACM CCS-8 Workshop DRM 2001, Philadelphia, PA, USA, November 5, 2001, Revised Papers*, volume 2320 of *Lecture Notes in Computer Science*, pages 76–105. Springer, 2001. doi: 10.1007/3-540-47870-1_6. URL https://doi.org/10.1007/3-540-47870-1_6.
- [46] Forbrukerrådet. Deceived by Design. URL <https://fil.forbrukerradet.no/wp-content/uploads/2018/06/2018-06-27-deceived-by-design-final.pdf>.
- [47] FreePrivacyPolicies.com. Free Privacy Policy Generator & Template with GDPR - Free Privacy Policy. URL <https://www.freeprivacypolicy.com/>.
- [48] Armin Gerl, Nadia Bennani, Harald Kosch, and Lionel Brunie. LPL, Towards a GDPR-Compliant Privacy Language: Formal Definition and Usage. 37:41–80.
- [49] GetTerms. Getterms.io. URL <http://getterms.io/>.
- [50] Google. Android Permissions overview, . URL <https://developer.android.com/guide/topics/permissions/overview>.
- [51] Google. Privacy Policy — Privacy & Terms — Google. Effective 22 January 2019, . URL <https://policies.google.com/privacy?gl=en&hl=en-GB>.
- [52] Colin M. Gray, Yubo Kou, Bryan Battles, Joseph Hoggatt, and Austin L. Toombs. The Dark (Patterns) Side of UX Design. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems - CHI '18*, pages 1–14. ACM Press. ISBN 978-1-4503-5620-6. doi: 10/gfxvpz. URL <http://dl.acm.org/citation.cfm?doid=3173574.3174108>.
- [53] Margaret D. Hagan. User-Centered Privacy Communication Design. URL <https://www.usenix.org/system/files/conference/soups2016/wfnp16-paper-hagan.pdf>.

- [54] Hamza Harkous, Kassem Fawaz, R mi Lebre t, Florian Schaub, Kang G. Shin, and Karl Aberer. Polisis: Automated Analysis and Presentation of Privacy Policies Using Deep Learning. URL <http://arxiv.org/abs/1802.02561>.
- [55] Manuel Hilty, Alexander Pretschner, David A. Basin, Christian Schaefer, and Thomas Walter. A Policy Language for Distributed Usage Control. In *Proceedings of the 12th European Symposium On Research in Computer Security, ESORICS'07*, volume 4734 of *Lecture Notes in Computer Science*, pages 531–546. Springer, 2007. ISBN 978-3-540-74834-2.
- [56] Gerard J. Holzmann. *The SPIN Model Checker - Primer and Reference Manual*. Addison-Wesley, 2004. ISBN 978-0-321-22862-8.
- [57] Michael Huth and Mark Dermot Ryan. *Logic in computer science - modelling and reasoning about systems (2. ed.)*. Cambridge University Press, 2004.
- [58] Iubenda. Features — Compliance Solutions, . URL <https://www.iubenda.com/en/features>.
- [59] Iubenda. Terms of service, . URL <https://www.iubenda.comhttps://www.iubenda.com/en/user/tos>.
- [60] Johnson Iyilade and Julita Vassileva. P2U: A Privacy Policy Specification Language for Secondary Data Sharing and Usage. pages 18–22. IEEE. ISBN 978-1-4799-5103-1. doi: 10.1109/SPW.2014.12. URL <http://ieeexplore.ieee.org/document/6957279/>.
- [61] Carlos Jensen and Colin Potts. Privacy policies as decision-making tools: An evaluation of online privacy notices. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 471–478. ACM. URL <http://dl.acm.org/citation.cfm?id=985752>.
- [62] Lalana Kagal. Rei. URL <http://ebiquity.umbc.edu/get/a/publication/57.pdf>.
- [63] Saffija Kasem-Madani and Michael Meier. Security and privacy policy languages: A survey, categorization and gap identification. URL <https://arxiv.org/abs/1512.00201>.
- [64] Patrick Gage Kelley, Joanna Bresee, Lorrie Faith Cranor, and Robert W. Reeder. A nutrition label for privacy. In *Proceedings of the 5th Symposium on Usable Privacy and Security*, page 4. ACM, . URL <http://dl.acm.org/citation.cfm?id=1572538>.
- [65] Patrick Gage Kelley, Lucian Cesca, Joanna Bresee, and Lorrie Faith Cranor. Standardizing privacy notices: An online study of the nutrition label approach. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 1573–1582. ACM, . URL <http://dl.acm.org/citation.cfm?id=1753561>.
- [66] Jan Kolter and G  nther Pernul. Generating User-Understandable Privacy Preferences. pages 299–306. IEEE. ISBN 978-1-4244-3572-2. doi: 10.1109/ARES.2009.89. URL <http://ieeexplore.ieee.org/document/5066486/>.
- [67] Ulrich K  nig and Jan Schallaboeck. Privacy preferences for E-Mail messages. URL <https://tools.ietf.org/html/koenig-privicons-03.txt>.
- [68] Marc Langheinrich, Lorrie Cranor, and Massimo Marchiori. Appel: A p3p preference exchange language. URL <https://www.w3.org/TR/P3P-preferences/>.

- [69] Daniel Le Métayer. A formal privacy management framework. In *International Workshop on Formal Aspects in Security and Trust*, pages 162–176. Springer. URL http://link.springer.com/chapter/10.1007/978-3-642-01465-9_11.
- [70] Andreas Matheus and J Herrmann. Geospatial Extensible Access Control Markup Language (GeoXACML).
- [71] Michael J. May, Carl A. Gunter, and Insup Lee. Privacy APIs: Access Control Techniques to Analyze and Verify Legal Privacy Policies. In *Proceedings of the 19th IEEE Computer Security Foundations Workshop, CSFW'06*, pages 85–97. IEEE Computer Society, 2006. ISBN 0-7695-2615-2.
- [72] Allison McCartney. How Lawyers Can Benefit From Visual Content. URL <https://visual.ly/blog/lawyers-visual-content/>.
- [73] Aleecia M. McDonald and Lorrie Faith Cranor. The cost of reading privacy policies. 4:543. URL http://heinonline.org/hol-cgi-bin/get_pdf.cgi?handle=hein.journals/isjlp soc4§ion=27.
- [74] Matthias Mehldau. Icons of privacy (original). URL <https://netzpolitik.org/wp-upload/data-privacy-icons-v01.pdf>.
- [75] Daniela Yidan Miao. PrivacyInformer: An Automated Privacy Description Generator for the MIT App Inventor. URL <http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.1029.2434&rep=rep1&type=pdf>.
- [76] V. Morel, M. Cunche, and D. Le Métayer. A generic information and consent framework for the iot. In *Proceedings of the 18th IEEE International Conference On Trust, Security And Privacy In Computing And Communications(TrustCom)*, pages 366–373, 2019. doi: 10.1109/TrustCom/BigDataSE.2019.00056.
- [77] Vivian Genaro Motti and Kelly Caine. Towards a Visual Vocabulary for Privacy Concepts. 60(1):1078–1082. ISSN 1541-9312. doi: 10/gf8vrk. URL <http://journals.sagepub.com/doi/10.1177/1541931213601249>.
- [78] National Telecommunications and Information Administration. Short Form Notice Code of Conduct to Promote Transparency in Mobile Apps Practices. URL https://www.ntia.doc.gov/files/ntia/publications/july_25_code-draft.pdf.
- [79] Netflix. Netflix Privacy Statement. Effective date: 11 May 2018. URL <https://help.netflix.com/en/legal/privacy>.
- [80] Helen Nissenbaum. Privacy as contextual integrity. 79:119. URL http://heinonline.org/hol-cgi-bin/get_pdf.cgi?handle=hein.journals/washlr79§ion=16.
- [81] Organisation for Economic Co-operation and Development. Skills matter: Further results from the survey of adult skills. OCLC: ocn953634518.
- [82] Raúl Pardo and Daniel Le Métayer. Analysis of privacy policies to enhance informed consent. In *Proceedings of the 33rd Annual IFIP WG 11.3 Conference on Data and Applications Security and Privacy*, volume 11559 of *Lecture Notes in Computer Science*, pages 177–198, 2019. doi: 10.1007/978-3-030-22479-0\10.

- [83] Jaehong Park and Ravi S. Sandhu. The UCON_{ABC} Usage Control Model. *ACM Trans. Inf. Syst. Secur.*, 7(1):128–174, 2004.
- [84] Niklas Paul, Welderufael B. Tesfay, Dennis-Kenji Kipker, Mattea Stelter, and Sebastian Pape. Assessing Privacy Policies of Internet of Things Services. In Lech Jan Janczewski and Mirosław Kutylowski, editors, *ICT Systems Security and Privacy Protection*, volume 529, pages 156–169. Springer International Publishing. ISBN 978-3-319-99827-5 978-3-319-99828-2. doi: 10.1007/978-3-319-99828-2_12. URL http://link.springer.com/10.1007/978-3-319-99828-2_12.
- [85] Polisis. Chrome Polisis, . URL <https://chrome.google.com/webstore/detail/polisis/bkddolgokpghlbhkhkflbbhhjghjdojck>.
- [86] Polisis. Firefox Polisis, . URL <https://addons.mozilla.org/en-US/firefox/addon/polisis/>.
- [87] Polisis. Polisis, . URL <https://www.priobot.org/polisis>.
- [88] Alexander Pretschner, Manuel Hilty, and David A. Basin. Distributed Usage Control. 49 (9):39–44. doi: 10/csvxdv.
- [89] Privacy Policy Generator. Privacy Policy Generator. URL <https://privacypolicygenerator.info/>.
- [90] Privacy Tech. Privacy icons. URL <https://www.privacytech.fr/privacy-icons/>.
- [91] PrivacyPolicies.com. Privacy Policy Generator: Free, GDPR, CalOPPA - PrivacyPolicies.com. URL <https://www.privacypolicies.com/>.
- [92] Aza Raskin. Making Privacy Policies not Suck, . URL <http://www.azaraskin.in/blog/post/making-privacy-policies-not-suck/>.
- [93] Aza Raskin. Privacy Icons - MozillaWiki, . URL <https://wiki.mozilla.org/Privacy-Icons>.
- [94] Joel R. Reidenberg, Jaspreet Bhatia, Travis Breaux, and Thomas B. Norton. Automated comparisons of ambiguity in privacy policies and the impact of regulation. . URL http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2715164.
- [95] Joel R. Reidenberg, Jaspreet Bhatia, Travis D. Breaux, and Thomas B. Norton. Ambiguity in Privacy Policies and the Impact of Regulation. 45(S2):S163–S190, . ISSN 0047-2530, 1537-5366. doi: 10/gdcdzm. URL <https://www.journals.uchicago.edu/doi/10.1086/688669>.
- [96] Joel R. Reidenberg, Travis Breaux, Lorrie Faith Cranor, Brian French, Amanda Grannis, James T. Graves, Fei Liu, Aleecia McDonald, Thomas B. Norton, and Rohan Ramanath. Disagreeable privacy policies: Mismatches between meaning and users’ understanding. 30:39, . URL http://heinonline.org/hol-cgi-bin/get_pdf.cgi?handle=hein.journals/berktech30§ion=6.
- [97] Arianna Rossi and Monica Palmirani. DaPIS: An Ontology-Based Data Protection Icon Set. pages 181–195. ISSN 0922-6389. doi: 10/gf7fbn. URL <http://www.medra.org/servlet/aliasResolver?alias=iospressISBN&isbn=978-1-61499-984-3&page=181&doi=10.3233/FAIA190020>.

- [98] Arianna Rossi, Rossana Ducato, Helena Haapio, and Stefania Passera. When Design Met Law: Design Patterns for Information Transparency. page 43.
- [99] Mary Rundle. International Data Protection and Digital Identity Management Tools, presentation at IGF 2006.
- [100] Norman Sadeh, Alessandro Acquisti, Travis D. Breaux, Lorrie Faith Cranor, Aleecia M. McDonald, Joel R. Reidenberg, Noah A. Smith, Fei Liu, N. Cameron Russell, Florian Schaub, et al. The usable privacy policy project. URL <http://ra.adm.cs.cmu.edu/anon/usr0/ftp/home/anon/isr2013/CMU-ISR-13-119.pdf>.
- [101] Ravi S. Sandhu, Edward J. Coyne, Hal L. Feinstein, and Charles E. Youman. Role-Based Access Control Models. 29(2):38–47. ISSN 0018-9162. doi: 10/bzxf36.
- [102] Florian Schaub, Rebecca Balebako, Adam L. Durity, and Lorrie Faith Cranor. A design space for effective privacy notices. In *Eleventh Symposium On Usable Privacy and Security (SOUPS 2015)*, pages 1–17. URL <https://www.usenix.org/conference/soups2015/proceedings/presentation/schaub>.
- [103] Bruce Schneier and John Kelsey. Cryptographic Support for Secure Logs on Untrusted Machines. page 11.
- [104] Meera Sivanathan. What is legal design? — Q&A with Meera Sivanathan (Legal Designer) — The Legal Forecast. URL <https://thelegalforecast.com/what-is-legal-design-qa-with-meera-sivanathan-legal-designer/>.
- [105] Daniel J Solove. Privacy Self-Management and the Consent Dilemma. page 25, .
- [106] Daniel J. Solove. A taxonomy of privacy. 154:477, . URL http://heinonline.org/hol-cgi-bin/get_pdf.cgi?handle=hein.journals/pnlr154§ion=20.
- [107] State of California. Assembly Bill No. 375 California Consumer Privacy Act.
- [108] Cass R. Sunstein. Choosing Not to Choose. ISSN 1556-5068. doi: 10/gftmr3. URL <http://www.ssrn.com/abstract=2377364>.
- [109] Latanya Sweeney, Mercè Crosas, and Michael Bar-Sinai. Sharing sensitive data with confidence: The datatags system. URL <http://techscience.org/a/2015101601/>.
- [110] Alasdair Taylor. Privacy policy. URL <https://seqlegal.com/free-legal-documents/privacy-policy>.
- [111] Terms of Service; Didn't Read. Terms of Service Classification. URL <https://tosdr.org/classification.html>.
- [112] Termsfeed. Generic Privacy Policy template. URL <https://www.termsfeed.com/assets/pdf/privacy-policy-template.pdf>.
- [113] Janice Y. Tsai, Serge Egelman, Lorrie Cranor, and Alessandro Acquisti. The Effect of Online Privacy Information on Purchasing Behavior: An Experimental Study. 22(2):254–268. ISSN 1047-7047, 1526-5536. doi: 10/cxhgzz. URL <http://pubsonline.informs.org/doi/abs/10.1287/isre.1090.0260>.
- [114] Twitter. Twitter Privacy Policy. Effective: May 25, 2018. URL https://cdn.cms-twdigitalassets.com/content/dam/legal-twitter/site-assets/privacy-page-gdpr/pdfs/PP.Q22018_April_EN.pdf.

- [115] United States Congress. Gramm—Leach—Bliley Act, .
- [116] United States Congress. Health Insurance Portability and Accountability Act, . URL <https://www.hhs.gov/sites/default/files/privacysummary.pdf>.
- [117] Christine Utz, Martin Degeling, Sascha Fahl, Florian Schaub, and Thorsten Holz. (Un)informed Consent: Studying GDPR Consent Notices in the Field. page 18.
- [118] Jasper van de Ven and Frank Dylla. Qualitative Privacy Description Language. In *Annual Privacy Forum*, pages 171–189. Springer. URL http://link.springer.com/chapter/10.1007/978-3-319-44760-5_11.
- [119] Bibi Van den Berg and Simone Van der Hof. What happens to my data? A novel approach to informing users of data processing practices. doi: 10.5210/fm.v17i7.4010. URL https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2100417.
- [120] Shomir Wilson, Florian Schaub, Aswarth Abhilash Dara, Sushain Cherivirala, Sebastian Zimmeck, Mads Scharup Andersen, Pedro Giovanni Leon, Eduard Hovy, and Norman Sadeh. Demystifying privacy policies with language technologies: Progress and challenges. .
- [121] Shomir Wilson, Florian Schaub, Aswarth Abhilash Dara, Frederick Liu, Sushain Cherivirala, Pedro Giovanni Leon, Mads Scharup Andersen, Sebastian Zimmeck, Kanthashree Mysore Sathyendra, N. Cameron Russell, et al. The creation and analysis of a website privacy policy corpus. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (ACL)*, . URL <http://www.aclweb.org/anthology/P/P16/P16-1126.pdf>.
- [122] WP29. Guidelines on transparency under Regulation 2016/679, .
- [123] WP29. Opinion 8/2014 on the Recent Developments on the Internet of Things. .
- [124] Jean Yang, Kuan Yessenov, and Armando Solar-Lezama. A language for automatically enforcing privacy policies. page 85. ACM Press. ISBN 978-1-4503-1083-3. doi: 10.1145/2103656.2103669. URL <http://dl.acm.org/citation.cfm?doid=2103656.2103669>.
- [125] Le Yu, Tao Zhang, Xiapu Luo, and Lei Xue. AutoPPG: Towards Automatic Generation of Privacy Policy for Android Applications. pages 39–50. ACM Press. ISBN 978-1-4503-3819-6. doi: 10.1145/2808117.2808125. URL <http://dl.acm.org/citation.cfm?doid=2808117.2808125>.
- [126] Zero-knowledge. Privacy Rights Markup Language Specification.
- [127] Sebastian Zimmeck and Steven M. Bellovin. Privee: An Architecture for Automatically Analyzing Web Privacy Policies. USENIX Association. ISBN 978-1-931971-15-7. URL <https://www.usenix.org/system/files/conference/usenixsecurity14/sec14-paper-zimmeck.pdf>. OCLC: 254320948.

A Glossary

To avoid any ambiguities, we highlight terms considered as important for the reading. As it is possible to find redundant terms in the literature, and different concepts expressed under the same word, we present a glossary to help the reader.

Personal data According to the GDPR, “personal data means any information relating to an identified or identifiable natural person (‘data subject’); an identifiable natural person is one who can be identified, directly or indirectly, in particular by reference to an identifier such as a name, an identification number, location data, an online identifier or to one or more factors specific to the physical, physiological, genetic, mental, economic, cultural or social identity of that natural person”.²³ Personal data can identify someone directly and uniquely — *e.g.* a social security number — with less precision — *e.g.* a pseudonym — or by combination with other information — *e.g.* metadata left by online behavior.

Data subject According to the GDPR, a data subject is “an identified or identifiable natural person”.²⁴ We name the data subject *DS* in this document.

Data controller According to the GDPR, “[data] controller means the natural or legal person, public authority, agency or other body which, alone or jointly with others, determines the purposes and means of the processing of personal data; where the purposes and means of such processing are determined by Union or Member State law, the controller or the specific criteria for its nomination may be provided for by Union or Member State law”.²⁵ We name the data controller *DC* in this document.

Privacy policy A privacy policy is a statement made by *DS* or *DC* to declare respectively their requirements and commitments in terms of personal data management. A privacy policy can be expressed in different ways: in natural language, graphically, or in machine-readable way.

Natural language privacy policies A *natural language privacy policy* is a privacy policy expressed in natural language.

Graphical privacy policies A *graphical privacy policy* is a privacy policy expressed graphically.

Machine-readable privacy policies A *machine-readable privacy policy* is a privacy policy expressed in a format readable by machines. This format is usually derived from a privacy policy language with a well-defined syntax and also in some cases a formal semantics.

Privacy policy language A *privacy policy language* is a language used to define privacy policies. It can describe *DC* as well as *DS* policies.

DC policy A *DC policy* is the privacy policy of a data controller. It is a commitment of the *DC* regarding its processing of personal data.

DS policy A *DS policy* is the privacy policy of a data subject. It defines the requirements of the *DS* concerning the processing of this data by *DC*.

Item An *item* is a piece of information provided in a privacy policy. Appendix B lists the items required by the GDPR.

²³See Article 4 of the GDPR.

²⁴See Article 4 of the GDPR, and the definition of *personal data* in this Glossary.

²⁵See Article 4 of the GDPR.

B Definition of policies

This section provides guidelines for the definition of *DS* and *DC policies* in line with the GDPR.

DC policy A *DC policy* should be understood as a privacy policy expressing the requests of DC related to the collection of DS personal data. It can be express in different ways, *i.e.* not only as a text but also as icons or in a formal language. Either way, the policy should at least provide the following items to answer the transparency requirements of the GDPR:

- the identity of the DC and its contact
- the type of data collected
- its purpose
- the legal basis for the processing
- the recipient of data
- the 3rd parties involved
- the retention time
- the rights of the DS

These requirements should be complemented with the following items to fully empower DS:

- the frequency of collection
- the location of the device
- its range of collection
- the beneficiary of the processing
- the risks associated
- what can be inferred from the collection and processing

On the one hand, location of device, frequency and range of collection are tailored to ubiquitous environments such as the IoT, and can provide better insights about devices. On the other hand, beneficiaries, risks and inferable information can summarize important information at a glance, and thereby provide more useful and impacting information.

DS policy A *DS policy* should be understood as a privacy policy expressing the requirements of DS related to the disclosure of their personal data. Among the settings DS should be able to adjust, the most obvious is their consent for data collection. And if the answer is positive, it has to be according to the following items:

- the type of data ²⁶
- the purpose of data collection
- 3rd party dissemination
- the retention time
- the data controller
- other requirements (anonymization, encryption, context ...)

²⁶Considering the difficulties a DS can have with technical terms, it should be possible to have a different granularity for the type of data, *e.g.* *unique identifiers* instead of MAC address and UID, geolocation etc



**RESEARCH CENTRE
GRENOBLE – RHÔNE-ALPES**

Inovallée
655 avenue de l'Europe Montbonnot
38334 Saint Ismier Cedex

Publisher
Inria
Domaine de Voluceau - Rocquencourt
BP 105 - 78153 Le Chesnay Cedex
inria.fr

ISSN 0249-6399