Detecting and Editing Privacy Policy Pitfalls on the Web

Cristiana Santos[1], Aldo Gangemi[2-3], Mehwish Alam[3]

[1] DH-CII, Universidade do Minho, Portugal
[2] ISTC–CNR, Italy
[3] LIPN, Université Paris 13, France

Abstract. Privacy policies are the locus where consequences concerning privacy and personal data are produced, but content features explain why they are largely ignored by its addressees. To abridge users with policies, we propose a policy-based system that identifies potential pitfalls in the privacy policies of companies on the Web. It will then suggest clarification of terms by suggesting removal or replacement of defective terms, in order to foster accountable policymaking and compliance. The proposed methods are based on extracting knowledge from natural language texts of a small sample size, and on semantic representations of the policy expression.

Keywords. Privacy policies, consumers, compliance, ontology, NLP, knowledge patterns, privacy and data protection

Introduction

Privacy policies consist of multiple paragraphs of natural language disclosing an organization's data practices on processing activities of personal data to its users, such as collection, use, sharing, and retention. They serve as a basis for decision-making [8], a "tool for preference-matching" for consumers [35], as consumers value a product/service more, after learning more about its attributes and tradeoffs for making a consumption decision. As such, they constitute the *locus*[1] where consequences are produced, the "*technically most feasible place to protect privacy and personal data*".

Notwithstanding its purposes and value, policy statements present concerns enumerated herein. There is no canonical format for presenting the information, and thus the language, organization, format and detail vary. Policy authors craft these policies broadly due to constant innovation, unstated present and future practices[2]. As explicitly stated by both consultants from Facebook and Google, "*Privacy policies are*

[1] President's Council of Advisors on Science and Technology: Big Data and Privacy: a Technological Perspective. Executive Office of the President, USA (2014), available online at https://bigdatawg.nist.gov/pdf/pcast_big_data_and_privacy_-_may_2014.pdf

[2] The motivations of policy authors for introducing vagueness in policy statements revolve from including unforeseeable situations, to accommodating flexibility by covering unknown and unstated existing internal practices. Surely, highly uncertain statements can easily accommodate a company's future practices, thus providing these companies more flexibility in the interim to alter those practices [15].

and will be written for lawyers and regulators because we are obliged to it"[3]. This cognitive overload over privacy policies will not change within the commercial realm. The use of unclear[4]-[5], open textured and ambiguous terms creates uncertainty, due to lack of information, or by leading to multiple interpretations. For instance, the fragment stating a broad purpose of "*improving customer experience*", or "*we disclose information to third parties only in aggregate or de-identified form*", and even "*we disclose certain personal information*", exemplifies vagueness in data practices, as it remains vague what information might be disclosed and which are the purposes. The use of complex language [6] [9] in the policies where instructions about opting out are obscured in the hinterlands constitutes also a problem. Moreover, privacy settings are also pointed down by its *form*, for they are generally concealed in small print and characterized by its lengthiness[6].

Content and form reasons elucidate why such statements are largely ignored and not used to change consumption decisions by its addressees[7] [3][19]. Despite all, privacy policies provisions are binding *inter partes*, regardless of whether or not their users read them, and only at court such clauses can be syndicated.

Contrastingly, researchers have also found that privacy and data protection infringements appear to be based also on developers' difficulties in understanding data protection and privacy requirements, rather than on malicious intentions [4]. We are observant of both strategic commercial practices, and also of the difficulties in understanding privacy policy statements by developers (mainly due to the porosity of policy language). We posit that privacy policies should be acquainted firstly by legal experts and knowledge engineers to inform any artefact.

Cognizant of the content features, in this paper we use the concept of privacy policy *pitfalls*. This concept conveys three requirements: i) it refers to commercial

[3] Declarations captured in a recent international workshop SC@Law - Fundamental Rights in the Digital Environments, held on the 19/05/2017 at the School of Law of the University of Minho).

[4] President's Council of Advisors on Science and Technology: Big Data and Privacy: a Technological Perspective. Executive Office of the President, USA (2014), available online https://bigdatawg.nist.gov/pdf/pcast_big_data_and_privacy_-_may_2014.pdf

[5] The Federal Trade Commission in the USA has found that most corporate privacy policies are 'incomprehensible' and that 'privacy policies do a poor job of informing consumers about companies' data practices or disclosing changes to their practices;' Preliminary Staff Report: Protecting Consumer Privacy in an Era of Rapid Change, March 2012. In March 2014, a French consumer group named UFC-Que Choisir**,** launched legal action against three of the largest social networks on grounds of breach of both consumer and data protection rules, having previously criticized the service providers for confusing ('*elliptique et pléthorique*') contractual terms; http://kiosque.quechoisir.org/magazine-mensuel-quechoisir-524-avril-2014/

[6] A study has calculated that it would take on average each internet user 244 hours per year to read the privacy policy belonging to each website they view, which is more than 50% of the time that average user spends on the internet [1].

[7] Few lay users (mostly the diligent ones) ever read privacy policies and regulators lack the resources to systematically review their contents or the ever emerging privacy policy notifications. According to the new Special Eurobarometer on data protection [3], only 18% of respondents fully read privacy statements, while 49% say they read them partially. Nearly a 31% say they don't read them at all. Indeed, standard privacy policies, broadly construed, do not make it easy for the average consumer to understand what is precisely made with the data collected about them. The survey explains that people who said that they do not fully read the privacy statements were asked to give their reasons for not doing so, and 67% of respondents say that they find the statements too long to read, while 38% find them unclear or too difficult to understand. Also, 15% of people say that they think the websites will not honour the statements anyway, 14% say they believe the law will protect them in any case, and 14% say that it is sufficient for them to see that websites have a privacy policy. 9% of respondents say they don't think it is important to read the privacy statements; 7% say that they don't know where to find them. The central finding of the survey shows that trust in digital environments remains low.

policy practices; ii) predicated in unclear, open textured and ambiguous terms; iii) previously syndicated by legal authorities and relevant stakeholders; this concept includes both linguistic (ii)) and legal elements (iii))[8].

It has been observed that pitfalls have seriously hampered privacy and personal data, because uncertain statements allow for interpretations that may be misleading, showing users a false sense of privacy [15]. Hence, the distinctive figure of this paper is tackling pitfalls of privacy policies in order to be understood by its lay recipients – the data subjects. The twofold objective of this paper is to describe some methods to i) identify potential pitfalls in the privacy policies of companies on the Web; and to ii) assist on the clarification of terms, for example, by replacing questionable terms of these documents. The proposed methods are based on extracting knowledge from natural language texts and semantic representations of the policy expression. The paper is organized as follows. Section 1 describes some examples of pitfalls in privacy policies. Details of the methods to identify and edit the pitfalls are given in Section 2; Section 3 describes related work, whereas Section 4 contains conclusions and future work description.

1. Privacy policy analysis: pitfall clauses showing non-compliance scenarios

This section presents the background on which we base our initial research. The content analysis of privacy policies has been grounded in authoritative sources, expert-generated documents, and correlated literature on privacy policy studies [8] [15] (cf. Section 2). Privacy policies apparently contain recurring textual frames that codify different data practices and enable identification of non-compliance scenarios reported therein. These patterns are consensuated and informed by the existing privacy and data protection framework in the EU[9].

We have further analysed data practice categories in order to formally represent *frames* (in the sense of [Fillmore [16], Gangemi [17]) that are typically associated with non-compliance scenarios emerging out of specific clauses.

We have singled out some clause patterns, which we call *pitfall clauses*, e.g. *advertising*, *amendment*, *connection clauses*. These patterns correspond to key user affordances (in the sense of [Gibson [16]]), e.g., *collection*, *retention*, *sharing*, *usage*.

We need to provide particular attention to the legal terms in a clause, compared against the domain legal knowledge.

i) "*Advertising clause*": with this clause, companies aim to use personal data (name, pictures, etc.) for advertising purposes. Despite the reference in the privacy settings, the

[8] First, we may be uncertain on whether a certain type of clause falls under the abstract legislative definition an "unfair contractual term". One can only have legal certainty that a certain type of clause is unfair if a competent institution, such as the European Court of Justice, has decided so.
In other cases the unfairness of a clause, has to be argued for, showing that it creates an unacceptable imbalance in the parties' rights and obligations. A consumer protection body might want to take the case to a court in order to authoritatively establish the unfairness of that clause, but a legal argument for that needs to be created, and the clause may eventually turn out to be judged fair.

[9] Including the GDPR, the Article 29 Working Party documents, EU Commission acts, Data Protection Authorities' enforcement actions, the decisions of the EU Court of Justice, amongst other documents.

context in which personal data may be used is not clear to its users (for example, if Facebook provides advertising spots to advertisers).
ii) "*Advertising clause: Information we receive – Information from other websites*": with this type of clause, companies wish to grant themselves the right to exchange data about its users with other websites, either affiliated or any other;
iii) "*Amendment clause": with this category clause, companies wish to reserve the right to make changes to the privacy policies at their sole discretion without the prior consent of the user. This means that substantial changes can be made without the users' knowledge. Such a clause is ineffective in its unrestricted formulation. Moreover, it contradicts two basic principles: the Lawfulness principle (the data subject has given his explicit consent for a specific purpose); and the Purpose Limitation principle (use of personal data for a purpose that is incompatible with the purpose(s) for which it was originally obtained). As an example, APPLE iCloud data privacy policies[10] do not promise to notify users about changes in the terms.
(iv) *Connection clause – "Information you share with others*": with this clause, companies wish to establish a connection to an application or website, by granting access to data (such as name, profile picture, gender, profession, etc.). This connection to apps and/or other websites is ineffective due to its lack of clarity and may violate the consent rule. The basis for evaluation should be the understanding of an "average consumer". According to this understanding, "access" to information on the Internet typically just means that third parties may view this information on the website. This clause, instead, relates to information that allows a data user to create profiles of the concerned individuals. This link or connection to an individual advertising profile created by the operator of the application or website clearly exceeds the consent given by the user from his or her point of view.

We illustrate the problem and motivate our approach using a running example. Let's account the WhatsApp case, denounced by the Article 29 Working Party (henceforth called WP29)[11]-[12]. The case concerns the *advertising clause* mentioned in i), known to each web service providers when they want to reuse user data. On the one hand, the excerpt WhatsApp clause[13] reads: "(...) Facebook and the other companies in the Facebook family also may use information from us to *improve your experiences* within their services such as *making product suggestions* (for example, of friends or connections, or of interesting content) and *showing relevant offers and ads* (...)". On the other hand, the WP29 denouncement states: "(...) WhatsApp will share information within the "Facebook family of companies" for a range of purposes that include marketing and advertising. *These are not purposes which were included within the Terms of Service and Privacy Policy when existing users signed-up to the service*. These changes have been introduced in contradiction with previous public statements of the two companies ensuring that no sharing of data would ever take place. (...) The

[10] Available at https://www.forbrukerradet.no/pressemelding/apple-icloud-violates-norwegian-and-european-law/

[11] The "Article 29 Working Party" is the short name of the Data Protection Working Party established by Article 29 of Directive 95/46/EC. It provides to the European Commission with independent advice on data protection matters and helps in the development of harmonised policies for data protection in the EU MS.

[12] Available at http://ec.europa.eu/justice/data-protection/article-29/documentation/other-document/files/2016/20161027__letter_of_the_chair_of_the_art_29_wp_whatsapp_en.pdf and the recent taskforce available at http://ec.europa.eu/newsroom/just/item-detail.cfm?item_id=50083

[13] This term is located under the epigraph "Affiliated Companies", in their website: https://www.whatsapp.com/legal/?l=en#privacy-policy-affiliated-companies

WP29 has serious concerns regarding the manner in which the information relating to the updated Terms of Service and Privacy Policy was provided to users and *consequently about the validity of the users' consent* (...)"(italics added). This denouncement is in line with the principles and rules of the General Data Protection Regulation[14] (hereafter called GDPR), as illustrated below:

- Lawfulness, fairness and transparency principles on processing personal data: "Personal data shall be processed lawfully, fairly and in a transparent manner in relation to the data subject ('lawfulness, fairness and transparency')", depicted in Article 5(1)(a);
- Purpose limitation principle on processing personal data: "Personal data shall be collected for specified, explicit and legitimate purposes and not further processed in a manner that is incompatible with those purposes ('purpose limitation')", depicted in Article 5(1)(b);
- Lawfulness of processing: "Processing shall be lawful only if and to the extent that at least one of the following applies: (a) the data subject has given consent to the processing of his or her personal data for one or more specific purposes", ascribed in Article 6(1)(a);

2. Methods to extract and represent pitfall clauses

Our ongoing research comprises a *bottom-up and practice-oriented* pipeline: document crawling, clustering, annotation, legal analysis, ontology design, automated knowledge extraction, and knowledge graph generation. A search of input documents related to privacy policies was performed. Mining these documents for potential pitfall terms and clauses was made. The legal analysis of the documents of the following types has enabled the identification of non-compliance scenarios, and their frame-based representation:

- Authoritative sources, e.g., the and issued judgments of the Court of Justice of the EU and national case-law[15];
- Expert-generated documents, such as the policy-based letters and opinions emanated by the stakeholders: Data Protection Authorities, the European Data Protection Supervisor, the WP29[16];
- Communications, complaints, written warnings, formal notices of correction, cease and desist letters from Consumer organizations[17] and Ombudsmen, reviewing and syndicating policy provisions or requesting their discontinuance, before filing lawsuits

Even if some of these documents are non-binding, their content may be relevant for pending or potential judicial procedures, and policy change. We produced a small sample size corpus of privacy policies of web companies whose business models are based on the commercialization of personal data, drawn from different categories (e.g.,

[14] All the articles cited in this paper refer to the GDPR.

[15] As an example, the Judgment by the Court of Appeal of 01/24/2014, Ref. No. 5 U 42/12 , Federation of German Consumer Organizations v. Facebook Ireland Limited.

[16] http://ec.europa.eu/newsroom/just/item-detail.cfm?item_id=50083

[17] E.g., the French UFC-Que Choisir, the European BEUC, and the Portuguese DECO have already summoned companies at court, among others.

entertainment, shopping, telecommunications, etc.). At this initial stage, we chose privacy policies of popular English-speaking websites that are visited by large numbers of users, specifically Whatsapp, Facebook, Apple and Microsoft. This analysis is required to check their consistency with the legal framework.

Terms and/or are considered pitfalls if they fulfil the mentioned conceptual requirements (open textured and ambiguous terms, and previously syndicated by legal authorities and relevant stakeholders) and are identified at least twice within the relevant documents. One can only have legal certainty that a term or clause is a pitfall if competent institutions, such as the European Court of Justice or a Data Protection Authority, have decided so. In other cases, the pitfall of a term/clause has to be argued for, showing that it creates an unacceptable imbalance between the parties' tradeoffs, and such legal arguments are delved by the domain stakeholders (known as data watchdogs). A development of a thorough taxonomy of pitfall is the first block to feed our pipeline and is being refined at the current moment. Table 1 shows a small fragment of the pitfalls typology.

Table 1. Fragment of the pitfalls typology

We disclose certain personal information
We disclose information to third parties only in aggregate or de-identified form
We collect your personal information, as necessary, to administer our business
WhatsApp is updating our Terms and Privacy Policy
Making product suggestions
Showing relevant offers and ads

The spotted policy frames will be applied to additional documents and refined[18] over multiple iterations against i) incoming front-runner documents from the concerned stakeholders on legal developments that affect one of the clauses denouncing non-compliance scenarios (recent court judgments, recommendations, complaints, notices of correction, cease and desist letters, etc); and ii) other privacy policies, until no further terms can be identified, thereby, assuring the saturation of terms and consolidating the process.

Contentious text fragments of explicit or implicit content portraying ambiguous terms are collected and annotated with their type, time stamp, and other metadata.

Each of the policy frames (or "knowledge patterns") of our typology, and its corresponding data compliance principles are used as a starting point to classify and annotate parts of a document, and derivatively the knowledge graphs extracted from them. For each pitfall, we propose alternative statements, either restricting the scope, or specifying purposes. For example, the broad purpose statement of collecting personal data for "*improving customer experience*", is evoked by the following clause: "*We collect your personal information in an effort to provide you with a superior customer experience and, as necessary, to administer our business*"[15]. We suggest that such collection should state *which is* the information collected and the *specific purposes* intended to improve customer service. Or, when we read "*WhatsApp is updating our*

[18] We assume that these frames are not static and others might be added to the initial categories.

Terms and Privacy Policy to reflect new features like WhatsApp calling", we further request *what* these new update *imply* (nature and implications).

Most of the initiatives that rely in NLP analysis of privacy policies try to code verbs and/or words for each of the data practices, e.g., the more exclusive verbs ''collect,'' ''disclose, and ''transfer''); however, among these keywords, a few words and phrases are not exclusively used to signal only one data practice, but many ("access,'' ''use", and' ''share", as the verb ''share'' indicates a transfer). Moreover, a policy may describe data collection, the purposes for collection, and data sharing requirements in different sections and under different titles, which can difficult our task of associating each title to each practice. For the purposes of our case study, we decided to exclude an analysis to the verbs themselves. Even if many linguistic categories indicate ambiguity in policies (conditionals, generalizations, modalities, quantifiers [15]), such categories are found in most of the clauses of policies, and this perspective would imply scrutinizing deeply each clause, without attending to other recurrent features that are contentious from both legal and linguistic perspectives that we have identified in our case study.

A legal ontology is aimed to describe the pitfalls, integrating existing ontological and non-ontological resources and adding domain-specific concepts.

Input documents are clustered, and their similarity measured by means of distributional techniques, i.e. by applying one or more distributional similarity measures for an efficient matching (e.g. WordNet-based [30], Explicit Semantic Analysis [31], word, sense, and frame embeddings [29]). We use FRED tool [10] [11] to extract frame-based knowledge graphs from the clustered documents, then we apply subgraph mining [11] to identify relevant patterns from multiple texts. Or readability purposes, an example is provided in Annex 1 with respect to the text of Article 6(1) (a) using FRED tool.

3. Related Work

There is substantial prior work in the area of expressing privacy policies. For the purposes of this paper we attend to two criteria: i) the domain of consumer privacy policies[19]; and ii) Machine-readable privacy statements, with a formal semantics based on an established formalism (RDF, OWL) so that we can more concretely foresee the implications of expressions in the language and consider the computational complexity of reasoning. We shall devote to the ones closest to our present work concerning consumer privacy policies expressed in semantic web technologies. We leave aside access control languages, and enterprise data flow requirements.

The Platform for privacy preferences (P3P) is an XML-based and privacy language for describing privacy practices of websites so that smart browsers could support consumers to check whether a policy conforms to a user's stated preferences [21]. Data requirement descriptions such as: no retention, purpose, legal requirement, business practices are important categories to consider. However, P3P cannot monitor compliance with the stated policy and it was declined as being too complex and

[19] The domain of privacy policies refers to the following typology: 1) online consumer privacy policies; (2 enterprise privacy policies, which govern an organization's internal business practices in relation to privacy; and 3) access control policies that implement a subset of enterprise privacy policy governing access to personal information;

confusing to be understood by an average user, and the P3P working group terminated their services in 2006.

Semantic web policy frameworks are based on OWL, used to express classification hierarchies with data type constraints for the semantic web. *KAoS* [22] is a framework that contains an OWL policy ontology, which consists of prescriptions as rights, prohibitions, obligations, and also exclusions. The OWL policy language *Rei* [23] provides classes for expressing rights, prohibitions, and obligations. Neither KAoS nor Rei ontologies detect conflicts between rights and prohibitions, nor infer rights from obligations; instead, to resolve modality conflicts, KAoS relies on a special priority-based algorithm, and *Rei* employs an RDF-S policy engine. *ExPDT* [24], another OWL-based policy language, focuses on conflict resolution via runtime monitoring. *MyCampus* [25] Project and *PeopleFinder* [26] project used a semantic web environment in which policies are expressed using a rule extension of the OWL language to: automate identification and access of personal resources, or capture privacy preferences, such as conditions under which users are willing to share their location or other user´s contextual resources with different services and other users. *Rein* is a semantic web framework [20] for representing and reasoning over policies in domains that use different policy languages and domain knowledge expressed in OWL and RDF-S, and supported rule languages (N3 rules and RuleML). It consists of two main parts, i) a set of ontologies for describing *Rein* policy networks and access requests; and ii) a reasoning engine that accepts requests for resources in *Rein* policy networks and decides whether or not the request is valid. *Rein* Ontology includes the *Rein* Policy Network Ontology, which describes the relationships between resources, policies, meta-policies, and policy languages, and the Request class, which is used to perform queries over *Rein* Policy Networks. This work is aimed at controlling access to resources and is domain independent.

As policy statements rely on language, relevant initiatives approached the ambiguity and vagueness of privacy policies on the web. In the work of Reidenberg, Bhatia, Breaux et al., [15] several ambiguous categories were considered: i) conditionals and conditional phrases (such as “when”, “upon”, and “during”); ii) generalizations (e.g. “typically” or “generally”); iii) modality (including modal verbs, like might, may, or, adverbs and non-specific adjectives); and iv) numeric quantifiers. These categories denote indeed vagueness, however, these are found in almost every clause. Scholars concluded that language has been crafted and specifically designed to give websites more flexibility, and as such, clauses are written in a more ambiguous way [15]. For example, the following clause contains all the detected ambiguous categories: conditions, generalizations, modal verbs and numeric quantifiers: “*we generally may share personal information we collect on the Site with certain service providers, some of whom may use the information for their own purposes as necessary*”. We are mostly centred in the current key frames we identified to spot the pitfalls, instead of scrutinizing each work in each clause.

The *Usable Privacy Policy* project 2016[20] [14] combines technologies, such as crowdsourcing, natural language processing (NLP), and machine learning to develop browser plug-in technologies that will automatically interpret privacy policies for users. The project extracts from a corpus of policies, in a semi-automated way, data practices, using crowdsourcing and NLP. This corpus of privacy policies was annotated by

[20] Usable Privacy Policy Project: https://www.usableprivacy.org/

experts with fine-grained detail about the data practices they contain; subtasks like identification of paragraph topics, user options will be automated. An analysis to the policies is achieved and then the policy features are translated into DL statements to facilitate detection of inconsistencies, contradictions, annotation disagreements, omissions and compliance violations. Preference modelling accounts privacy concerns, perceptions, preferences, cognitive biases that may negatively affect individuals' privacy decisions. Finally, the project encompasses a user interface for privacy notices (e.g. labels and icons). Its bottom-up approach and fine-grained analysis served as an inspiration to our model, but this project holds only for the USA jurisdiction, and we perceive that the new GDPR data rules (EU-wide) will have a major impact on companies of all sizes worldwide. Moreover, this Regulation offers a robust principle-based framework towards privacy and data protection.

Ontologies applied to privacy policies have been developed to support privacy-preserving systems. Yet, there is a visible preterition of ontology-based systems able to analyse policy pitfalls, verifying whether policy content actually complies with the GDPR, and reporting its results to organizations; formal models of computational ontologies representing the GDPR are few and thoroughly designed considering theoretical aspects, ultimately difficult to be used in practical settings, such as the *Ontology to Model Data Protection Requirements in Workflows* [12] which deals with the GDPR, encoding rights and obligations of both data controllers and data subjects, but leaves aside policy segments, the centre of our study. *PrivOnto* [13], built upon the *Usable Privacy Policy* project, consists in a semantic technology framework that models, in a machine-readable format, US-based privacy policies, in order to answer to privacy questions of interest to users. In particular, it is aimed at: retrieving salient statements (ambiguous, inconsistent and incomplete) made in privacy policies, modeling their contents using ontology-based representations; and using semantic web technologies to explore the obtained knowledge structures. The ontology populated starting from the annotation schema and the corpus (populated with about 23,000 annotations of data practices). This ontology is still in a development stage and not assessed yet by the scientific community.

Eddy [27] consists in a DL language (interlingua). This syntax was designed to describe privacy requirements specifications in order to automate conflict detection of privacy policies, enabling the alignment of data flows (chains) from third-party services or platforms, and across multi-tier applications. This formal language was applied to real-world policies from Facebook, Zynga, and AOL Advertising. In this approach, natural language requirements statements are mapped to "actions" and "roles" in DL to check consistency and detecting requirements conflicts within single party's privacy specification, and conflicts between two or more specifications of different parties. "Actions" correspond to usage, transfer, collection and retention; and each action definition is expressed using the "roles": hasObject, hasSource, hasPurpose, and hasTarget. The six role concepts for the TBox include: "modality" (whether the action is a permission, obligation, prohibition), "actor" (the actor who performs the action on the datum), "datum" (the information on which the action is performed), "purpose" (the purpose for which the action is performed); source (the source from which the information is collected), target (for transfer actions, the recipient to whom the information is transferred); these concept are organized into a hierarchy using DL subsumption, as it is suitable for expressing and reasoning over ambiguity that frequently appears in natural language requirements. The process of coding the policy statements implies that the analyst annotates and parameterizes each text-based data requirement and assigns one of the four action codes (transfer, use, collection, retention) into their roles and role values. Then the analysts use 3 additional codes to extract subsumption relationships: role abstractions, refinements, and exclusions which are often used by policy writers[21] to illustrate by example the types of information that are acted upon. The analyst transfers the encoded values into the Eddy language that employs an SQL-like syntax and the DL semantics. Breaux et al. [28] later extended this representation with notions of rights, obligations and permissions in a case study, and then formalized this extension in DL.

Tools allow downloading privacy policy templates[22], through a privacy policy generator, contributing to the lengthy and legalese born-digital clauses.

Advances over the state of the art in NLP information extraction reveal tools such as *IBM Watson*[23]. But *IBM Watson* is not based on deep parsing, has no advanced semantic role labeling, and does not produce formal graphs, as it is the case with *FRED*. Legal AI applications like IBM's Ross [32] are still obscure and skepticism about its performance has been drawn by [33] [34].

4. Conclusions and Future Work

This paper has described the theoretical background of an ongoing research that uses semantic web techniques to automatically identify pitfalls in data clauses which are subject to the GDPR rules. It aims to identify privacy-policy pitfalls and to propose a possible clarification of terms.

We observed that even though privacy policies communicate data-handling practices, they are crafted with an ambiguous wording.

The underlined methods consisted so far in the annotation and extraction of the knowledge patterns found in natural language texts.

Such an approach, even in a small set of policies, is still labor-intensive and withstands with the complexity and updating of language of privacy policies. However, the encoding of pitfall policies will have impact at several stances. By describing privacy policy pitfalls in a machine-readable representation format, may enable interoperability and reusability by organizations that oversee data protection issues. Also, by making use of semantic technologies, data will come with rich descriptions within its context but connected to other entities in the web of data. Our endeavor can help businesses in devising and drafting lawful privacy policies and achieving better transparency, trust, competitive advantages, which would also benefit users. The methodology and technical tools might enable regulators to easily spot poor privacy policies and empower them to more effectively target enforcement actions. Clarity in privacy practices is a necessary prerequisite for empowering users to make informed decisions about upholding their data. This system is part of the privacy by design principle, as a proactive and preventative measure.

[21] 1) Role abstractions, which consist of one or more concepts that are more generic than a given role value (e.g., ‘‘information’’ is more generic than ‘‘a person’s name’’); 2) role refinements, which consist of one or more concepts that are more specific than a given role value (e.g., ‘‘browser type’’ and ‘‘screen resolution’’ are specific kinds of ‘‘technical information’’); and 3) role exclusions, which consist of one or more concepts that are excluded from a given role value (e.g., ‘‘IP addresses’’ are excluded from what a company might consider ‘‘personal information’’).

[22] https://termsfeed.com/privacy-policy/generator/

[23] https://www.ibm.com/watson/

In the future, we intend to improve the coverage of policies of the present study. We aim at verifying empirically if an edited policy clause makes it easy for users to limit the ways in which the company collects their personal information, i.e., if policies add value to decision-making, in particular as a tool for preference-matching. We envision that contracts, agreements and other high value documents that organizations may possess can be further analyzed with our system to ensure compliance within the new regulatory environment. The ineffectiveness of some clauses will also consider domestic legal requirement from national privacy laws from France, Italy and Spain to ensure legal compliance at national level and not only within the range of the Regulation. Privacy policies can be also compared against two official benchmarks to show whether official privacy disclosures result in policies less ambiguous than those edited by our system. Moreover, we wish to check which data practices (collection, sharing, usage, retention) correspond to more pitfall clauses. We aim to contribute to an EU Model Privacy Form, as the template used in the USA, in the financial domain, under the Gramm-Leach-Bliley Act[24]. In our study, we don´t consider complex composition of services, where policies of other platforms and third-party data flows are combined, but such ecosystem will be considered in the near future.

References

[1] McDonald A. M., Cranor L. F. (2008), The cost of Reading Privacy Policies, A Journal of Law and Policy for the Information Society 2008, Privacy Year in Review, p.17.

[2] Preliminary Opinion of the European Data Protection Supervisor Privacy and competitiveness in the age of big data: The interplay between data protection, competition law and consumer protection in the Digital Economy, March 2014, available online at https://edps.europa.eu/sites/edp/files/publication/14-03-26_competitition_law_big_data_en.pdf

[3] Special Eurobarometer 431 - Data protection, June 2015

[4] Balebako R., Marsh A., Lin J., et al. (2014), The privacy and security behaviors of smartphone app developers, in USEC, 2014.

[5] Reidenberg, J.R., Bhatia, J., Breaux, T.D., et al. (2017) Automated Comparisons of Ambiguity in Privacy Policies and the Impact of Regulation. J Legal Studies 47 (forthcoming).

[6] Reidenberg, J.R., Breaux, T.D., Cranor, L.F., et al. (2015). Disagreeable Privacy Policies: Mismatches between Meaning and Users' Understanding. Berkeley Technology Law Journal 39, 1, 2015.

[7] Hoke, C., Cranor, L.F., Leon, P.G., et al. (2014) Are They Worth Reading? An In-Depth Analysis of Online Advertising Companies' Privacy Policies, 2014 TPRC Conference Paper.

[8] Jensen, C., Potts, C. (2004). Privacy policies as decision-making tools: an evaluation of online privacy notices. Proc. SIGCHI conference on Human factors in computing systems (CHI), 471– 578.

[9] Fader A., Soderland S., Etzioni O. (2011): Identifying Relations for Open Information Extraction. EMNLP.

[10] Gangemi A., Reforgiato D., Mongiovì D. et al. (2016). Identifying motifs for evaluating open knowledge extraction on the Web, Knowledge-Based Systems (108), 33-41

[11] Gangemi A., Presutti V., Reforgiato D. et al. (2016). Semantic Web Machine Reading with FRED, Semantic Web Journal

[12] Bartolini, C. Muthuri R. et al. (2015). Using Ontologies to Model Data Protection Requirements in Workflows, Ninth International Workshop on Juris-informatics (JURISIN)

[13] Oltramari A., Piraviperumal D., Schaub F., et al. (2016), PrivOnto: a Semantic Framework for the Analysis of Privacy Policies, Semantic Web Journal (under review)

[14] Wilson S., Schaub F. , Dara A.A. et al, (2016), The creation and analysis of a website privacy policy corpus, Proc. 54th Annual Meeting of the Association for Computational Linguistics (ACL), 2016

[24] https://www.ftc.gov/sites/default/files/documents/rules/privacy-consumer-financial-information-financial-privacy-rule/model_form_rule_a_small_entity_compliance_guide.pdf

[15] Reidenberg J., Bhatia J., Breaux T., et al, (2016). Ambiguity in Privacy Policies and the Impact of Regulation, 45 Journal of Legal Studies

[16] Fillmore C. J. (1976) Frame semantics and the nature of language. Annals of the New York Academy of Sciences, 280(1):20–32

[17] Gangemi A. (2010). What's in a Schema? Cambridge University Press, Cambridge, UK, pp. 144–182.

[18] Gibson, J. J. (1977). The theory of affordances. In R. Shaw & J. Bransford (Eds.), Perceiving, acting, and knowing: Toward an ecological psychology (pp. 67-82). Hillsdale, NJ: Erlbaum

[19] Calo, M R. (2012) Against notice skepticism in privacy (and elsewhere), Notre Dame Law Review, 87, pp. 1027-2261

[20] Kagal, L., Berners-Lee, T., Connolly, D. et al. (1999) Using semantic web technologies for policy management on the web. Proceeding of the National Conference on Artificial Intelligence, Vol. 21. No. 2. Menlo Park, CA; Cambridge, MA; London; AAAI Press; MIT Press

[21] Cranor L et al (2006) Platform for privacy preferences (P3P) specification. W3C working group note

[22] Uszok A, Bradshaw JM, Lott J, Breedy M, Bunch L (2008), New developments in ontology-based policy management: increasing the practicality and comprehensiveness of KAoS. In: IEEE workshop on policies for distributed systems and networks, pp 145–152

[23] Tonti G, Bradshaw JM, Jeffers R, et al. (2003) Semantic web languages for policy representation and reasoning: a comparison of KAoS, Rei, and Ponder. LNCS 2870:419–437

[24] Kahmer M, Gilliot M, Muller G (2008) Automating privacy compliance with ExPDT. In: 10th IEEE conference on e-commerce technology, pp 87–94

[25] Sadeh, N., Gandon, F., and Oh Buyng, K. (2006) Ambient Intelligence: The MyCampus Experience. In Ambient Intelligence and Pervasive Computing. Eds. T. Vasilakos andW. Pedrycz, ArTech House

[26] Gandon, F.L., and Sadeh, N.M.: Semantic web technologies to reconcile privacy and context awareness. Web Semantics: Science, Services and Agents on the World Wide Web. 241-260 (2004)

[27] Breaux, T.D., Hibshi, H., and Rao, A. (2014), Eddy, a formal language for specifying and analyzing data flow specifications for conflicting privacy requirements. Requirements Engineering, 19.3, 281-307

[28] Breaux TD, Anton AI (2005) Analyzing goal semantics for rights, permissions, and obligations. In: IEEE international requirements engineering conference, Paris, France, pp 177–186

[29] Alam M., Recupero D.R., Mongiovì M. et al. (2017),Event-based knowledge reconciliation using frame embeddings and frame similarity. Knowl.-Based Syst. 135: 192-203 (2017)

[30] Christiane Fellbaum (1998) WordNet: an electronic lexical database. MIT Press, 1998.

[31] Gabrilovich E., Markovitch S.,(2007). Computing semantic relatedness using Wikipedia-based explicit semantic analysis. In IJCAI'07: Proc. 20th International Joint Conference on Artificial Intelligence.

[32] Beck S. (2014) The Future of Law The Americal Lawyer, 2014

[33] Paliwala A. (2016) Rediscovering artificial intelligence and law: an inadequate jurisprudence? International Review of Law, Computers & Technology 30:107-114

[34] Remus D, Levy FS (2015) Can Robots Be Lawyers? Computers, Lawyers, and the Practice of Law

[35] Bambauer, J. (2017) Are Privacy Policies Information or Ideological?, 66 DePaul L. Rev.

Annex 1: Article 6(1) (a)
is captured by FRED tool.

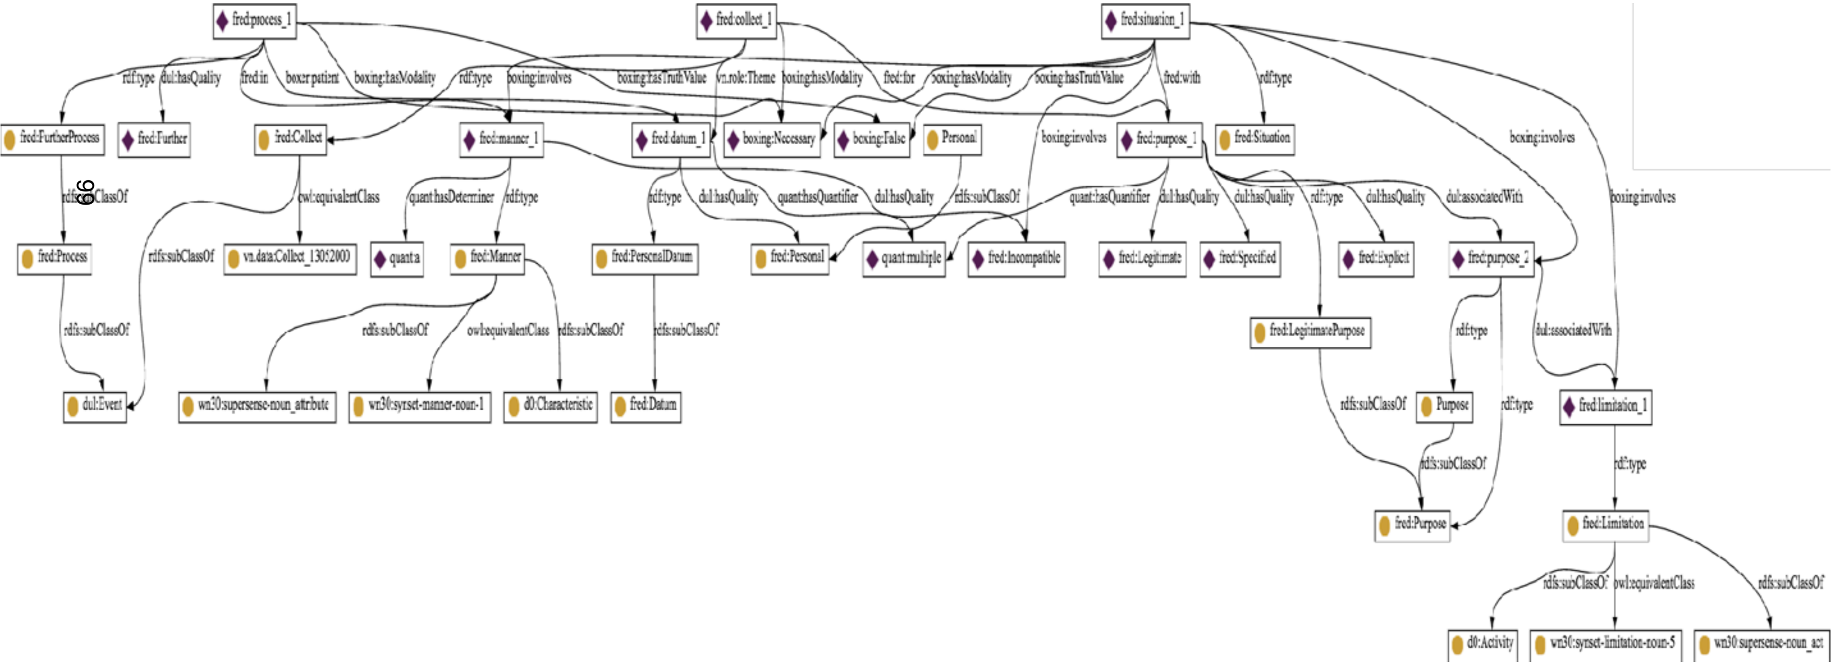
fred:process_1
fred:collect_1
fred:situation_1
rdf:type
dul:hasQuality
fred:in
boxer:patient
boxing:hasModality
rdf:type
boxing:involves
boxing:hasTruthValue
vn.role:Theme
boxing:hasModality
fred:for
boxing:hasModality
boxing:hasTruthValue
fred:with
rdf:type
boxing:involves
boxing:involves
fred:FurtherProcess
fred:Further
fred:Collect
fred:manner_1
fred:datum_1
boxing:Necessary
boxing:False
Personal
fred:purpose_1
fred:Situation
rdfs:subClassOf
owl:equivalentClass
quant:hasDeterminer
rdf:type
rdf:type
dul:hasQuality
quant:hasQuantifier
dul:hasQuality
rdfs:subClassOf
quant:hasQuantifier
dul:hasQuality
dul:hasQuality
rdf:type
dul:hasQuality
dul:associatedWith
fred:Process
rdfs:subClassOf
vn.data:Collect_13052000
quant:a
fred:Manner
fred:PersonalDatum
fred:Personal
quant:multiple
fred:Incompatible
fred:Legitimate
fred:Specified
fred:Explicit
fred:purpose_2
rdfs:subClassOf
rdfs:subClassOf
owl:equivalentClass
rdfs:subClassOf
rdfs:subClassOf
fred:LegitimatePurpose
rdf:type
dul:associatedWith
dul:Event
wn30:supersense-noun_attribute
wn30:synset-manner-noun-1
d0:Characteristic
fred:Datum
rdfs:subClassOf
Purpose
rdf:type
fred:limitation_1
rdfs:subClassOf
rdf:type
fred:Purpose
fred:Limitation
rdfs:subClassOf
owl:equivalentClass
rdfs:subClassOf
d0:Activity
wn30:synset-limitation-noun-5
wn30:supersense-noun_act