

A Semantic Tool for the Protection of Personal Information Act

Yahlieel Jafta

2858132@myuwc.ac.za

University of the Western Cape

Cape Town, Western Cape

ABSTRACT

The increase of connectivity due to technology advances is shifting the attention of legislators in various jurisdictions to the protection of personal information and data. The focus of this paper is the protection of privacy information, specifically the POPIA within South Africa. It is yet to be implemented and role-players within the POPIA still have time to be compliant and understand how it affects them. This paper follows the development of an ontology to establish a small knowledge base for the regulations in the POPIA and how it affects these role players. It provides a background on semantic technologies and outlines the scope and the development of an ontology.

KEYWORDS

Legal ontology, data protection, information privacy, Protection of Personal Information Act, General Data Protection Regulation

1 INTRODUCTION

We are living in an era in which connectivity is progressing rapidly. During the past decade, this increase in connectivity shifted the focus of legislators in various jurisdictions to the protection of personal information. The motivation for this focus is a matter of concern, due to the impact this has on the privacy rights of individuals.

Embedded in section 14 of the Constitution of the Republic of South Africa, is the right to privacy by each South African citizen. This includes individuals and businesses. “The right to privacy includes a right to protection against the unlawful collection, retention, dissemination, and use of personal information” [10]. While businesses have justifiable reasons to acquire personal data as information assets to reach their business goals, they are required to comply with any legislation involving the processing of such data [3]. Nevertheless, this information is susceptible to abuse. With the rise of social media, businesses have new means to gain traction in the information space. A prime example of this was brought to the fore on the international stage by the recent “Facebook–Cambridge Analytica data scandal” [8]. It was uncovered that Cambridge Analytica procured the personal information of countless individuals’ Facebook profiles in the absence of their consent and used it for political purposes, incurring gross breaches of privacy [8]. In the wake of this, the European Union (EU) was changing its present data protection law, the General Data Protection Regulation (GDPR). This change was created since 2010 and became effective on 25 May 2018. Therefore, legislators are in the process of enacting laws to secure proper information handling procedures.

The objective for domains of law with regards to data protection and privacy is to ensure the security of individuals’ personal information in jurisdictions [4]. Organizations operational in South

Africa are confronting the authorization of the Protection of Personal Information Act 4 of 2013 (POPIA) which has been signed into law by the President on 19 November 2013 [37]. The enforcement of the POPIA is yet to commence, and the regulations to the Act were only recently published on 14 December 2018. Once implementation is confirmed, there will be a one year grace period for entities to be compliant, before enforcement. The Act aims to align the regulation of personal information in South Africa with international standards, which will put pressure on businesses operating within South Africa.

The POPIA seeks to protect the right of privacy that applies to individuals and juristic entities (referred to as data subjects) by the establishment of strict guidelines on how to obtain and process information [37]. These guidelines affect organizations (referred to as the responsible party [37]) and data subjects in various ways and are therefore important that these role players are well informed with regards to the implications. It is with this in mind that a knowledge base, through a semantic representation, for legislation of the POPIA will be valuable for assisting with the education of both the data subjects and organizations on the POPIA.

In this paper, the author provides the analysis, design, and implementation of a basic ontology on the data protection domain with regards to the POPIA. The paper is arranged as follows. Section 2 provides a background on semantic technologies, their importance, and influences in the scope of knowledge bases and information formalization. Section 3 describes related work in regards to legal domain ontologies concerning the privacy and protection of data. Section 4 will outline the functional and non-functional requirements. The methodology for developing the ontology will be discussed in section 5. Section 6 outlines the design and implementation of the ontology and present a prototype. Section 7 discusses the initial evaluation of the ontology followed by a summary of the paper and future development.

2 BACKGROUND

2.1 Semantic Technologies

One of the biggest problems we face today is an overload of information. This is evident in various domains as the availability of large scale information is more abundant than ever before. “Businesses that operate globally have become more dependent on information systems to survive” [23]. This is amplified by the huge size of the World Wide Web (WWW) and the resources it provides. In the WWW context, search engines made amazing progress in handling vast amounts of information and making it available on the Web. However, with the continuous growth of the Web, search engines will find it hard to identify relevant results for search terms [5]. An argument is made for Semantics being the single most important factor for advancing the Web to its next phase.

Semantic technology encodes meanings separately from data. It deals with large data sets and links them together via self-describing interrelations, allowing it to be processed by machines. Semantic technologies are one of the initiatives to address the challenge that large scale information, known as Big Data, presents. The advances in hardware along with Big Data provides enormous opportunities for individuals, businesses, and society, but we haven't seen similar advances in software development capability. This could potentially give rise to a second Software Crisis [16]. Semantic technology seeks to address this potential crisis in the Big Data space. It is viewed as the leading framework to deal with the diverse and huge size of assets on the Web [35]. Even though the Web is a major force in the scope of Big Data, it is not the only push factor. Push factors include the processing of large data in enterprises and various domains, where making good decisions efficiently is essential to allowing organizations to manage and make better sense of their data. The Semantic Web is an expansion of the present Web, wherein data is given unambiguous meaning [35]. This enables machines and people to coordinate various actions [6].

2.2 Ontology

In Computer Science, “an ontology is a formal representation of the knowledge by a set of concepts within a domain and the relationships between those concepts” [26]. It is the process of formalizing knowledge and expressing the concepts and their relations in a given domain. “An ontology defines a common vocabulary for researchers who need to share information in a domain” [13]. It incorporates machine–interpretable meanings of fundamental ideas within the domain and the relationships between them [13]. As a result, ontologies introduces a sharable and reusable knowledge base, enabling the extension of knowledge of a given domain. Ontologies are a big factor in promoting automated processes to access information. It is relied upon to produce organized vocabularies that investigate the connections between various terms, enabling machines and people to decipher their “meaning” unambiguously [20]. For example, “a pizza ontology might include the information that Mozzarella and Parmesan are variants of cheese, that cheese is not a kind of meat or seafood, and that a vegetarian pizza is one whose toppings do not include any meat or seafood. This information allows the term *pizza topped with only Mozzarella and Parmesan* to be unambiguously defined as a specialization of the term *vegetarian pizza*” [20]. The formal representation of the information in an ontology allows for the extension of existing knowledge. This is achieved by inferences made on the existing knowledge base. Inference is a tool to improve the quality of data by discovering new relationships and performing automated analyses on the content of the data to extend the existing knowledge or identify any data inconsistencies.

Ontologies are used in various domains as a type of information depiction about the world or some subset of it [26]. Domains include Artificial Intelligence, Semantic Web, Biomedical Informatics and even in the legal domain. The use of ontologies in the legal domain assists with the organization of legal documents and providing support for legal reasoning [3]. Even though the focus of this paper is on the legal domain we have seen several successful ontologies

developed in the health sector domains. This includes the Gene Ontology [9] and the Protein ontology [36].

2.2.1 Gene Ontology (GO). Applying to the domain of biology. GO is a bioinformatics asset that represents the responsibility of genes in organisms, covering a variety of species from humans to bacteria and viruses [29]. “The GO project seeks to provide a set of structured vocabularies for specific biological domains that can be used to describe gene products in any organism” [9]. The goal was to develop a knowledge base of terms relating to various organisms, which are then consumed by several databases to explain genes and gene products unambiguously [9].

2.2.2 Protein ontology (PO). The motivation for the development of the PO was to “efficiently represent the protein annotation framework and integrating the existing data representations into a standardized protein data specification for the bioinformatics community” [36]. The objective of the PO was “to correlate information about multiprotein machines with data in major protein databases to better understand the sequence, structure, and function of protein machines” [36].

2.3 OWL

The Web Ontology Language (OWL) [27] is “a Semantic Web language intended to represent and capture rich and complex knowledge about things, and their relations between them” [34]. The knowledge expressed by OWL enables it to be reasoned about. This is achieved utilizing automated reasoners that verify the consistency of the domain knowledge within an ontology or revealing hidden knowledge [34]. OWL ontologies promote reuse and modularity, as it can be published in the WWW and be referenced from other OWL ontologies [34]. Knowledge concepts captured from data, in a given domain, are reasoned about in a rich hierarchical structure of concepts and their inter-relationships [36]. These relationships help with the matching of concepts even if the data sources describing these concepts are not 100% uniform.

OWL is not limited to the Web as there have been successful applications of knowledge modeling in several application areas, such as [9] and the [36]. Modeling knowledge in OWL has two focal points. “As a descriptive language, it can be used to formalize domain knowledge, and as a logical language, it can be utilized to make inferences from this knowledge” [24]. The latter point sets OWL apart from other modeling languages, for example, UML [24].

Since 2009 there is a second version of OWL. This is due to the challenges faced in the initial version, OWL 1. The challenges presented relates to the efficiency and scalability of the reasoning process. The reasoning was typically a multi-exponential process, which did not always yield results [24]. To address these issues, OWL 2 [2] was released. OWL 2, a World Wide Web Consortium (W3C) standard, presented three profiles: “OWL EL, OWL RL, and OWL QL” [24]. Also called sublanguages [24]. These sublanguages restrict accessible modeling features to streamline reasoning. This has yielded great results in improving the efficiency and scalability which made the OWL 2 profiles appealing for ontology engineers [24].

“OWL EL is used in large biomedical ontologies” [24], with applications of this in [9] and [36]. OWL RL is the favored method for

reasoning upon information on the WWW [24]. OWL QL enhances information access of applications integrating with databases. This is demonstrated in [21].

The usage of OWL is applied in various successful ontologies, thus demonstrating the benefits it provides.

3 RELATED WORK

The POPIA is very closely aligned with the GDPR [12], with the latter being more expressive. However, one of the differences is that the POPIA applies the term data subject to natural and juristic persons, whereas the GDPR only applies this to natural persons [12]. There are various approaches related to expressing jurisdictional regulations, such as the GDPR, as an ontology. The development of an ontology is inspired by these ontologies, but based on the POPIA. As part of the literature review, a thorough search of the relevant literature on the POPIA yielded no related articles. The objective of this paper is the utilization of an ontology to demonstrate concepts expressed within the POPIA and how it influences the data subjects and responsible parties once the Act is implemented. The goal is to create a knowledge base, as a proof of concept, for the POPIA concerning the lawful processing of personal information. The source for the descriptions of the POPIA concepts will be derived from the official text described within the Act [37] documentation. Below, the relevant domain ontologies are described highlighting their implementation and goal.

Pandit et al. [32] developed an ontology, GDPR text extensions (GDPRtEXT), that provides an approach to allude to the ideas and terms conveyed inside the GDPR. It was developed using the “Simple Knowledge Organization System” (SKOS) [28] as a source for the descriptions of the GDPR concepts. “SKOS is a Semantic Web language for representing formally structured vocabularies” [28]. The terms are linked to the appropriate focuses in the GDPR text using a URI pattern which links each term to a distinct resource within the GDPR. The GDPRtEXT ontology does not make use of inference to provide a better understanding of compliance obligations.

An ontology for the GDPR was developed for the data protection requirements [4]. It shows the data protection prerequisites with regards to the GDPR change and introduces a methodology for incorporating it into a work process to express these necessities inside a business procedure through the ontology. The goal of the ontology is to provide support for data controllers in accomplishing consistency with the GDPR enactment [4]. This was done to create an ontological representation of the obligations of data controllers, and the comparing privileges of information subjects [4].

The GDPRov project [33] is an ontology concerned with the management of compliance through recognizing provenance information identified with assent and individual personal information required for consistency documentation. “It is an OWL 2 linked open data ontology” [33] that represents the provenance of assent and data lifecycle work processes for the GDPR. It outlines the provenance of exercises, for example, “data securing, usage, storage, deletion, and sharing of consent and the life cycles of data” [33]. The ontology uses SPARQL [7] to query the provenance information described to find information relevant for compliance.

The PrOnto ontology [31] is a privacy ontology that conceptualizes the main concepts in the GDPR. These include “data types

and documents, agents and roles, processing purposes, legal bases, processing operations, and deontic operations for modeling rights and duties” [31]. The objective of PrOnto is to assist with legal thinking and verification of compliance [31]. It achieves this by applying defeasible logic reasoning as opposed to solely information retrieval. “Defeasible reasoning is a rule-based method for efficient reasoning with inconsistent data” [22].

4 REQUIREMENTS

4.1 Functional Requirements

The requirements for this project will be defined by a set of competency questions the ontology should answer. These questions will serve as the litmus test in the evaluation phase of the development process and will help define the scope of the ontology [13]. An initial question list is defined below.

- What is the responsibility of the various role-players?
- Who is responsible for ensuring the conditions for processing data comply with the Act?
- What is considered personal information?
- How do organizations dealing with international organizations in the management of data comply with the Act?
- How does POPIA prevent the transfer of data across South African borders?
- How does POPIA allow the transfer of data across South African borders?

The set of questions is still in the process of evaluation and is subject to change during evaluation and maintenance.

4.2 Non-functional Requirements

The development tool that will be used for this project is Protégé [30]. Protégé is a free open-source ontology editor that was created at the “Stanford Center for Biomedical Informatics Research at the Stanford University School of Medicine”. It is upheld by an enormous network of scholarly, government, and corporate clients, who use it to create knowledge-based solutions for different domains [17]. One of the main strengths of protégé is its user interface and the flexible manner in which it can be extended to provide additional functionality in the form of plug-ins [25]. One of the design goals of Protégé is to be compatible and adaptable with other systems for knowledge representation and knowledge extraction [25]. This compatibility will be a great feature for a POPIA ontology as the information privacy domain is large, enabling possible integration with other knowledge bases within the same domain.

The ontology language that will be used for this project is OWL and it is available as a plugin for protégé. The OWL plugin allows users to make full use of the features of OWL within the protégé editor enabling them to make use of an intelligent development environment comparable to a software programming environment. The features and benefits outlined for the combination of protégé and the OWL plugin align with the non-functional requirements that need to be satisfied. Since the ontology is based on an Act, it will always be susceptible to change, due to changes or amendments in regulations. Based on this, an initial list of non-functional requirements that will need to be satisfied is Adaptability, Reusability,

Configurability, Testability, Maintainability, and Quality. Considering these requirements, the protégé editor provides a suitable environment for the development of an ontology to meet these requirements.

5 METHODOLOGY

This project will be implemented through four phases: Analysis, Design, Implementation and Evaluation. Analysis: This involves a literature study and review. Defining the requirements for the project within the scope of the POPIA. Design: Includes the modeling of various concepts within the POPIA. Implementation: The implementation forms part of the design phase, as the modeling is done within the protégé editor. This allows for quick feedback in the design phase. The evaluation will be performed at the latter stages of the project to test the conformance of ontology to requirements.

A progression of strategies for creating ontologies has been accounted for in literature since the mid-1990's demonstrating that there is no single set of processes for ontology engineering. Methodologies include The Cyc methodology, methodology of Uschold and King, methodology of Grüninger and Fox, METHONTOLOGY, The KACTUS approach, and the SENSUS-based methodology [14].

The Cyc methodology is a product of the development of the Cyc Knowledge Base (Cyc KB) [11]. Each phase includes the development of knowledge representation, outlining abstract concepts and representing the rest of the knowledge using these concepts [14].

The methodology defined by Uschold and King [38] is based on the development of the Enterprise Ontology, “an ontology for enterprise modeling processes at the Artificial Intelligence Applications Institute (AIAI) of Edinburgh” [38].

The methodology of Grüninger and Fox was established through the development of the “Toronto Virtual Enterprise (TOVE) project” [18]. “An ontology in the domain of business processes” [19]. It involves building a logical model of a knowledge source that is to be specified using an ontology. It is a “logic-based formal methodology” that formalizes informal scenarios, expressed in a natural language, into a machine-readable model that can be processed [14].

The METHONTOLOGY method enables the development of ontologies at the knowledge level and follows an iterative approach utilizing evolving prototypes. The establishment of this methodology is in the primary exercises distinguished by the software development process [1]. It supports prototyping and comprises of the following processes: “Specification, Conceptualization, Formalization, Implementation, and Maintenance”.

The “Ontology Development 101: A Guide to Creating Your First Ontology” [13] presents another method for the development of ontologies. It is an iterative development process that repeats continuously to enhance the ontology. It consists of the following sub-processes:

- “Determine the domain and scope by defining a set of competency questions.
- Explore the reuse of existing ontologies.
- Listing key terms in the ontology.
- Create the classes and class hierarchy.
- Create the properties of classes.
- Create features for the defined properties.

- Create instances” [13].

The methodology framework that will be used for building an ontology-based on the POPIA is a combination of the METHONTOLOGY [15] and the Ontology Development 101 [13] method. These two methods have complementary processes to assist with development. The METHONTOLOGY provides high-level activities for the development life cycle and the ontology development 101 methods provide granular steps for design and implementation, which are intricate phases in the development life cycle. Furthermore, both these methodologies follow an iterative development model which is satisfactory for the development of this project.

6 DESIGN AND IMPLEMENTATION

The design follows the sub-processes of the Ontology Development 101 guide, specified in section 5. It consists of the following processes: Listing all the key terms considered important for the knowledge base, defining an initial class hierarchy, defining properties for classes and their features. To reduce scope creep, the reuse of existing ontologies will not be considered for this version of the POPIA ontology. It will be a consideration for future work.

6.1 Important Terms

Describing the list of terms follows the process of identifying terms from the Act that will assist in explaining concepts in the POPIA. The terms are chosen based on the competency questions the ontology should answer. To enable reasoning about the terms, the properties of the terms are also taken into consideration.

The glossary of terms includes, but not limited to, *data subject, processing, accountability, right to access, operator, person, responsible party and personal information*.

For determining the classes, a list of key terms from the glossary was identified that describe the concepts having independent existence [13]. The list of terms was then grouped according to the main concepts they represent within the Act. These groups form the basis of the ontology architecture. It is made up of the following main concepts:

- The Conditions for lawful processing of personal information
- The rights of data subjects
- Data Processing
- Person
- Action

These concepts are used to conceptualize and establish the class definitions of the ontology. The following sub-sections describe the derivation of the classes and properties.

6.2 Classes

6.2.1 The Conditions for lawful processing of personal information. Chapter 3 of the Act defines the conditions for the lawful processing of personal information. It details eight conditions, with sub-conditions, that are to be complied with by the Responsible Party. “The responsible party must ensure that the conditions set out in this Chapter, and all the measures that give effect to such conditions, are complied with at the time of the determination of

Table 1: Core class Hierarchy

Base Classes	Subclasses
LawfulCondition	Accountability, InformationQuality, Lawfulness, Openness, PurposeSpecification, Security
DataSubjectRight	RightToAccess, RightToCorrection, RightToDelete, RightToObject, RightToSubmitComplaint
DataProcessing	LawfulProcessing, ProcessingActivity, ProcessingMode
Person	DataSubject, Child, ResponsibleParty, Operator, Regulator, CompetentPerson
Action	DataSubjectAction, ResponsiblePartyAction

Table 2: Core properties

Property	Domain	Range
exerciseRight	DataSubject	DataSubjectRight
hasInformation	DataSubject	PersonalInformation
mustEnsure	ResponsibleParty	LawfulCondition
performProcessing	Operator	DataProcessing

the purpose and means of the processing and during the processing itself” [37].

The list of conditions is “Accountability, Processing limitation, Purpose specification, Further processing limitation, Information quality, Openness, Security Safeguards and Data subject participation” [37].

A subset of these concepts was used to define the classes for the lawful processing of personal information. From this list, the classes created were the following: Accountability, InformationQuality, Lawfulness, Openness, PurposeSpecification, Security. These classes are subclasses of the LawfullCondition root class.

6.2.2 The rights of data subjects. Chapter 2 Section 5 of the Act defines the rights of data subjects. “A data subject has the right to have his, her or its personal information processed following the conditions for the lawful processing of personal information as referred to in Chapter 3” [37]. The concepts used for describing these rights as classes are derived from the nine rights described. These rights include, but not limited to, “the right to be notified when personal information is collected, the right to request correction of personal information and the right to object to the processing of personal information” [37].

From these rights, a DataSubjectRight class was created to represent these rights with subclasses representing the individual rights. The subclasses include RightToAccess, RightToCorrection, RightToDelete, RightToObject and RightToSubmitComplaint

6.2.3 Data processing. Chapter 1 provides the definitions and purpose of the Act. In this chapter personal information processing is defined as follows. “Processing means any operation or activity or any set of operations, whether or not by automatic means, concerning personal information” [37]. This includes the collection, storage, transfer or destruction of information to list a few definitions. More detail is provided in Chapter 1 of the Act. Additionally, Chapter 8 Section 71 provides additional information on personal information processing. “A data subject may not be subject to a decision which results in legal consequences for him, her or it, or which affects him, her or it to a substantial degree, which is based solely on the basis of the automated processing of personal information

intended to provide a profile of such person including his or her performance at work, or his, her or its credit worthiness, reliability, location, health, personal preferences or conduct” [37]. This section provides additional details on processing and highlights automated processing and profiling which are used to further describe data processing.

These concepts are used to extend the class definitions for conceptualizing data processing. These include DataProcessing with subclasses ProcessingActivity, AutomatedProcessing, InformationTransfer, Profiling, ManualProcessing

6.2.4 Person. In Chapter 1 of the Act, a person is defined as “a natural person or a juristic person” [37]. Based on this, it describes the following concepts as a person. A child, who is “a natural person under the age of 18 years who is not legally competent” [37]. A competent person is any “person who is legally competent to consent to any action or decision being taken in respect of any matter concerning a child” [37]. A data subject refers to “the person to whom the personal information relates” [37]. An operator is “a person who processes personal information for a responsible party” [37]. The “responsible party is a public or private body or any other person which, determines the purpose of and means for processing personal information” [37]. Finally, the last concept identified for defining the “person” concept is the regulator. Chapter 5 Section 39 defines the regulator as a juristic person.

From these definitions, the following classes are created: Person, Regulator, Child, DataSubject, CompetentPerson, Operator and ResponsibleParty.

6.2.5 Action. There exists a link between the responsible party and a data subject due to the lawful conditions required for information processing and the rights that data subjects have. This link was used to describe two new concepts, a data subject action, and a responsible party action. A data subject has the right to exercise their rights such as the right to objection of personal information processing, Chapter 2 Section 5 of the Act. One of the conditions of lawful information processing is the responsibility of the responsible party to notify data subjects when their personal information is accessed and collected.

These two concepts form the basis for concepts described as an Action, since it is an action that has to be performed by the responsible party, and in the case of the data subject, can be performed.

Table 1 describes the core classes.

6.3 Properties

The properties are created by identifying the terms that provide the correlation between the classes defined and then expressing these correlations as relations within the ontology.

A subset of terms is listed below that was used to form the core properties of the ontology:

- “A data subject can **exercise his/her rights**.”
- A data subject has **personal information**.
- Data processing is **performed by** an operator.
- The responsible party to **ensure** conditions for lawful processing.” [37]

The highlighted terms are used to create properties. For example, exercise rights relate a data subject to an action they can perform, which in turn is related to a data subject right. An operator performs data processing for a responsible party, thus creating a relationship between the Operator, DataProcessing and ResponsibleParty classes.

Table 2 describes the core properties.

6.4 Concepts and Relations from related work

There is some overlap between the regulations of the GDPR and the POPIA. Thus some concepts from the ontology developed by Bartolini et al. [4] were used, and in some instances changed, to fit the POPIA ontology. The following high-level concepts were used.

The GDPR ontology [4] has a concept “Action” which represents the actions of data subjects. This concept was used in the POPIA ontology as a root class and represented the actions of data subjects and the responsible party as a DataSubjectAction and ResponsiblePartyAction class. The following properties was used: *accessData*, *exerciseRight*, *objectTo*, *performProcessing* and *processingPerformedBy* [4].

7 EVALUATION

An initial evaluation was performed on the ontology by performing some basic queries using the DL Query plugin in the Protégé editor. The queries included querying the rights of data subjects, describing personal information, special personal information and creating an initial set of instances. The evaluation is still in progress and will be continued in the next phase.

8 CONCLUSION AND FUTURE WORK

This paper highlights the importance of private information regulation in a globally connected world. It highlights the shift in the focus of legislators to enact legislation to support the privacy rights of individuals and juristic entities. The focus of this paper is the South African POPIA and the impact it will have on various role players, such as the responsible party and data subjects. It outlines the development of an ontology to provide a knowledge base on various concepts within the Act that will promote transparency

and education that can aid with the inception of this Act. The development of the ontology is still in progress. A set of classes and properties was created to demonstrate a working ontology. The ontology will be the input for the next development phase, evaluation. The evaluation of the ontology will be performed to assert the satisfaction of requirements followed by maintenance. Initial evaluation was performed on the ontology and a few classes were identified that requires refactoring to enable requirement satisfaction. This will form part of the next phase, followed by maintenance.

REFERENCES

- [1] 1996. IEEE Standard for Developing Software Life Cycle Processes. *IEEE Std 1074-1995* (1996), i-. <https://doi.org/10.1109/IEEESTD.1996.79663>
- [2] 2012. *OWL 2 Web Ontology Language Document Overview (Second Edition)*. W3C Recommendation. W3C. <http://www.w3.org/TR/2012/REC-owl2-overview-20121211/>.
- [3] Cesare Bartolini. 2015. Reconciling Data Protection Rights and Obligations : An Ontology of the Forthcoming EU Regulation.
- [4] Cesare Bartolini, Robert Muthuri, and Cristiana Santos. 2017. Using Ontologies to Model Data Protection Requirements in Workflows. In *New Frontiers in Artificial Intelligence*, Mihoko Otake, Setsuya Kurahashi, Yuiko Ota, Ken Satoh, and Daisuke Bekki (Eds.). Springer International Publishing, 233–248.
- [5] V. Richard Benjamins, Jesus Contreras, Oscar Corcho, and Asuncion Gomez-Perez. 2002. Six challenges for the Semantic Web. (04 2002).
- [6] Tim Berners-Lee. [n. d.]. The Semantic Web. <https://www.scientificamerican.com/article/the-semantic-web/>
- [7] Kendall Grant Clark, Lee Feigenbaum, and Elias Torres. 2008. SPARQL Protocol for RDF. World Wide Web Consortium, Recommendation REC-rdf-sparql-protocol-20080115.
- [8] MacKenzie F. Common. 2018. Facebook and Cambridge Analytica: let this be the high-water mark for impunity. <http://eprints.lse.ac.uk/88964/>.
- [9] T. G. O. Consortium. 2001. Creating the Gene Ontology Resource: Design and Implementation. *Genome Research* 11, 8 (Aug 2001), 1425–1433. <https://doi.org/10.1101/gr.180801>
- [10] Constitution of the Republic of South Africa. 1996. Section 14. <http://www.justice.gov.za/legislation/constitution/SACConstitution-web-eng.pdf>.
- [11] Cycorp. [n. d.]. Knowledge Base. <http://www.cyc.com/kb/>
- [12] Adéle Da Veiga, Ruthea Vorster, Fudong Li, Nathan Clarke, and Steven Furnell. 2018. A comparison of compliance with data privacy requirements in two countries. In *26th European Conference on Information Systems*. University of Portsmouth.
- [13] N F. Noy and Deborah McGuinness. 2001. Ontology Development 101: A Guide to Creating Your First Ontology. *Knowledge Systems Laboratory* 32 (01 2001).
- [14] Mariano Fernández-López and Asunción Gómez-Pérez. 2002. Overview and Analysis of Methodologies for Building Ontologies. *Knowl. Eng. Rev.* 17, 2 (June 2002), 129–156. <https://doi.org/10.1017/S0269888902000462>
- [15] M. Fernández-López, A. Gómez-Pérez, and N. Juristo. 1997. METHONTOLOGY: From Ontological Art Towards Ontological Engineering. In *Proceedings of the Ontological Engineering AAAI-97 Spring Symposium Series*. American Association for Artificial Intelligence. <http://oa.upm.es/5484/> Ontology Engineering Group ? OEG.
- [16] Brian Fitzgerald. 2012. Software Crisis 2.0. *IEEE Computer - COMPUTER* 45 (04 2012), 89–91. <https://doi.org/10.1109/MC.2012.147>
- [17] Stanford Center for Biomedical Informatics Research. [n. d.]. A free, open-source ontology editor and framework for building intelligent systems. <https://protege.stanford.edu/> [Online; accessed 24-March-2019].
- [18] Mark S. Fox. 1992. The TOVE Project Towards a Common-Sense Model of the Enterprise. In *IEA/AIE*.
- [19] Michael GrÄijninger and Mark S. Fox. 1995. Methodology for the Design and Evaluation of Ontologies.
- [20] Ian Horrocks, Peter Patel-Schneider, and Frank Harmelen. 2003. From SHIQ and RDF to OWL: The Making of a Web Ontology Language. *SSRN Electronic Journal* (01 2003). <https://doi.org/10.2139/ssrn.3199003>
- [21] Evgeny Kharlamov, Nina Solomakhina, Özgür Lütfü Özçep, Dmitriy Zheleznyakov, Thomas Hubauer, Steffen Lamparter, Mikhail Roshchin, Ahmet Soylu, and Stuart Watson. 2014. How Semantic Technologies Can Enhance Data Access at Siemens Energy. In *The Semantic Web – ISWC 2014*, Peter Mika, Tania Tudorache, Abraham Bernstein, Chris Welty, Craig Knoblock, Denny Vrandečić, Paul Groth, Natasha Noy, Krzysztof Janowicz, and Carole Goble (Eds.). Springer International Publishing, 601–619.

A Semantic Tool for the Protection of Personal Information Act

- [22] Efstratios Kontopoulos, Nick Bassiliades, Guido Governatori, and Grigoris Antoniou. 2013. A Modal Defeasible Reasoner of Deontic Logic for the Semantic Web. *Semantic Web* (2013), 140–167. <https://doi.org/10.4018/978-1-4666-3610-1.ch007>
- [23] Selçuk Köylüoğlu, Levent Duman, and Aykut Bedük. 2015. Information Systems in Globalization Process and Their Reflections in Education. *Procedia – Social and Behavioral Sciences* 191 (06 2015), 1349–1354. <https://doi.org/10.1016/j.sbspro.2015.04.610>
- [24] Markus Krötzsch. 2012. *OWL 2 Profiles: An Introduction to Lightweight Ontology Languages*. Springer Berlin Heidelberg, Berlin, Heidelberg, 112–183. https://doi.org/10.1007/978-3-642-33158-9_4
- [25] Patrick Lambrix, Manal Habbouche, and Marta Pérez. 2003. Evaluation of ontology development tools for bioinformatics. *Bioinformatics* 19 12 (2003), 1564–71.
- [26] Diana Man. 2013. Ontologies in Computer Science. *Didactica Mathematica* 31 (2013), 43–46.
- [27] Deborah McGuinness and Frank van Harmelen. 2004. *OWL Web Ontology Language Overview*. W3C Recommendation. W3C. <http://www.w3.org/TR/2004/REC-owl-features-20040210/>.
- [28] Alistair Miles and Josiah R. Pérez-Agüijera. 2007. SKOS: Simple Knowledge Organisation for the Web. *Cataloging & Classification Quarterly* 43, 3-4 (2007), 69–83. https://doi.org/10.1300/J104v43n03_04 arXiv:https://doi.org/10.1300/J104v43n03_04
- [29] Christopher J Mungall, Heiko Dietze, and David Osumi-Sutherland. 2014. Use of OWL within the Gene Ontology. In *OWLED*.
- [30] Mark A. Musen. 2015. The protégé project: a look back and a look forward. *AI Matters* 1, 4 (2015), 4–12. <https://doi.org/10.1145/2757001.2757003>
- [31] Monica Palmirani, Michele Martoni, Arianna Rossi, Cesare Bartolini, and Livio Robaldo. 2018. *PrOnto: Privacy Ontology for Legal Reasoning: 7th International Conference, EGOVIS 2018, Regensburg, Germany, September 3–5, 2018, Proceedings*. 139–152. https://doi.org/10.1007/978-3-319-98349-3_11
- [32] Harshvardhan J. Pandit, Kaniz Fatema, Declan O’Sullivan, and Dave Lewis. 2018. GDPRtEXT - GDPR as a Linked Data Resource. In *ESWC*.
- [33] Harshvardhan J. Pandit and Dave Lewis. 2017. Modelling Provenance for GDPR Compliance using Linked Open Data Vocabularies. In *PrivOn@ISWC*.
- [34] Bijan Parsia, Peter Patel-Schneider, Markus Krötzsch, Sebastian Rudolph, and Pascal Hitzler. 2012. *OWL 2 Web Ontology Language Primer (Second Edition)*. Technical Report. W3C. <http://www.w3.org/TR/2012/REC-owl2-primer-20121211/>.
- [35] Amit Sheth and Cartic Ramakrishnan. 2004. Semantic (Web) Technology In Action: Ontology Driven Information Systems For Search, Integration and Analysis. *IEEE Data Engineering Bulletin* 26 (01 2004).
- [36] Amandeep Sidhu, Tharam S. Dillon, Elizabeth Chang, and Baldev S. Sidhu. 2005. Protein Ontology Development using OWL., Vol. 188.
- [37] South African Government Gazette. 2013. Protection of Personal Information Act. <http://www.justice.gov.za/legislation/acts/2013-004.pdf>.
- [38] Michael Uschold and Martin King Aiai. 1995. Towards a Methodology for Building Ontologies.

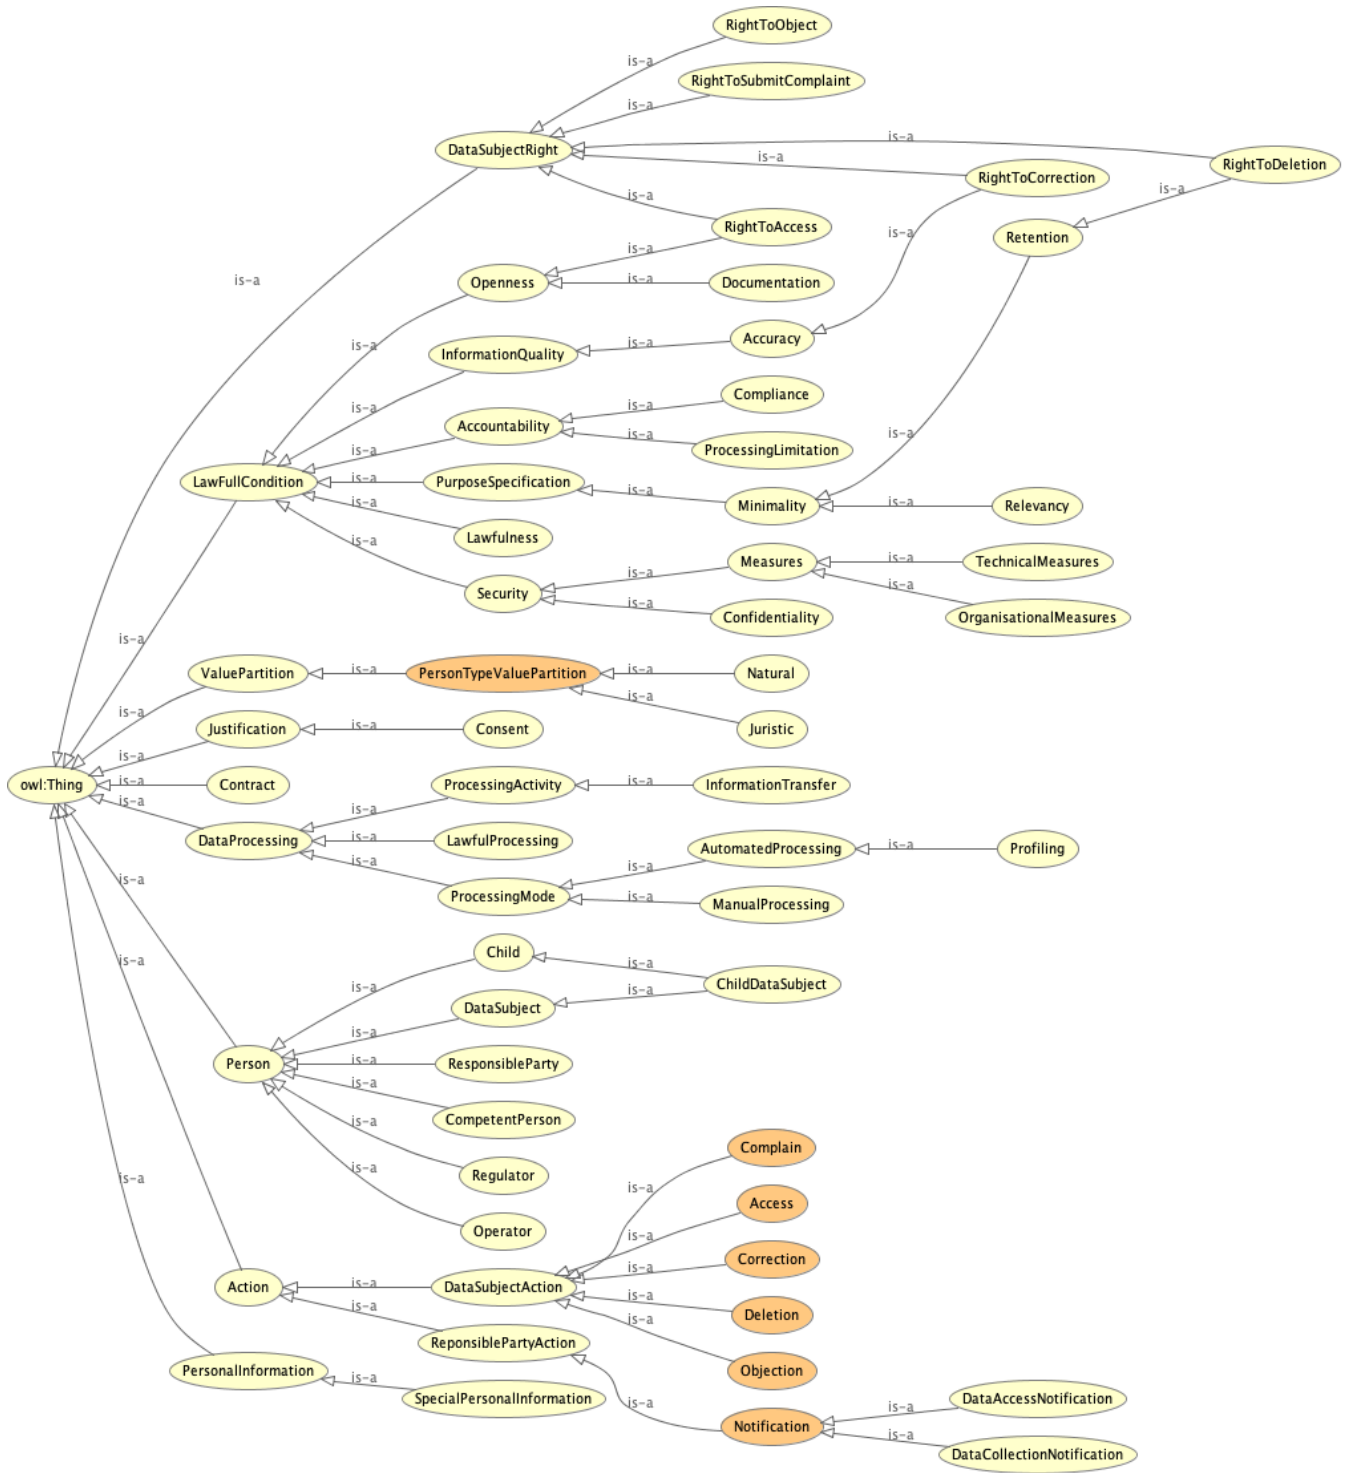


Figure 1: Ontology schema