

AN ONTOLOGY FOR THE SOUTH AFRICAN PROTECTION OF PERSONAL INFORMATION ACT

Y. Jafta¹, L. Leenen^{1,2}, and P. Chan³

¹ University of the Western Cape, South Africa

² CAIR, South Africa

³ Council for Scientific and Industrial Research (CSIR), Pretoria, South Africa

[1 2858132@myuwc.ac.za](mailto:12858132@myuwc.ac.za)

[1,2 lleenen@uwc.ac.za](mailto:1,2lleenen@uwc.ac.za)

[3 kchan@csir.co.za](mailto:3kchan@csir.co.za)

Abstract: The protection and management of data, and especially personal information, is becoming an issue of critical importance in both the business environment and in general society. Various institutions have justifiable reasons to gather the personal information of individuals but they are required to comply with any legislation involving the processing of such data. Organisations thus face legal and other repercussions should personal information be breached or treated negligently. Most countries have adopted privacy and data protection laws or are in the process of enacting such laws. In South Africa, the Protection of Personal Information Act (POPIA) was formally adopted in 2013 but it is yet to be implemented. When the implementation of the Act is announced, role players (responsible parties and data subjects) affected by POPIA will have a grace period of a year to become compliant and/or understand how the Act will affect them. One example of a mandate that follows from POPIA is data breach notification. This paper presents the development of a prototype ontology on POPIA to promote transparency and education of affected data subjects and organisations including government departments. The ontology provides a semantic representation of a knowledge base for the regulations in the POPIA and how it affects these role players. The POPIA is closely aligned with the European Union's General Data Protection Regulation (GDPR), and the POPIA ontology is inspired by similar ontologies developed for the GDPR.

Keywords: Legal Ontology, Data Protection, Information Privacy, Protection of Personal Information Act, General Data Protection Regulation

1. Introduction

The rapid increase in connectivity and gathering of data have shifted the focus of legislators in various jurisdictions to the protection of personal information. This shift in focus has an impact on the privacy rights of individuals. Most countries have adopted, or are in the process of, adopting privacy laws. The objective of legislation with regards to data protection and privacy is to ensure the security of individuals' personal information in jurisdictions (Bartolini et al., 2015b). The European Union (EU) introduced the Data Protection Directive (DPD) as early as 1995 (Birnhack, 2008). The EU revised this directive into the General Data Protection Regulation (GDPR) in 2015 and updated it in 2018 on the reformed it (Srncova et al., 2019). The United Kingdom adopted the Data Protection Act in 1998 (United Kingdom Government Gazette, 1998) in addition to the EU DPD and this was implemented in 2000 (Botha et al., 2017). The United States of America does not have a single comprehensive law regulating the use and collection of personal information but has enacted several privacy laws since 2001 (Information Shield, nd). In South Africa, the Protection of Personal Information (POPI) Act was enacted on November 26th, 2013 (South African Government Gazette, 2013) although the full enforcement date is yet to be determined by the country's privacy regulator. South African businesses, government departments, and civil society is in the process of complying with the Act, but are facing implementation challenges and lack of awareness of the impact of the POPI Act. The Act aims to align the regulation of personal information in South Africa with international standards and the South African constitution. Embedded in section 14 of the Constitution of the Republic of South Africa, is the right to privacy of each South African citizen. The right to privacy includes a right to protection against the unlawful collection, retention, dissemination, and use of personal information" (Constitution of the Republic of South Africa, 1996).

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 (Bartolini & Mathuri, 2015a). Nevertheless, such 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 the recent "Facebook Cambridge Analytica data scandal". 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 (Common, 2018). In the wake of this, the EU reformed its data protection law, the General Data Protection Regulation (GDPR). These changes became effective on 25 May 2018 (Srncova et al., 2019).

Whilst enforcement of the South African POPI Act 4 of 2013 (POPIA) is yet to commence, the regulations to the Act were only published on 14 December 2018 (South African Government Gazette, 14 December 2018). Once implementation is confirmed, there will be a one-year grace period for entities to become compliant. 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 (South African Government Gazette, 2013). These guidelines affect organisations (referred to as the responsible party) and data subjects in various ways, and it is 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. The POPIA is closely aligned with the European Union's General Data Protection Regulation (GDPR), and the POPIA ontology is inspired by similar ontologies developed for the GDPR (Pandit et al., 2018, Miles et al., 2007).

This paper provides the analysis, design, and implementation of a basic ontology on the data protection domain with regards to the POPIA. The purpose of the POPIA ontology is to support understanding and awareness of the Act, and for future versions to be integrated into systems where automated compliance checking may be required. Section 2 provides a background on semantic technologies, their importance, and influences in the scope of knowledge bases and information formalisation. Section 3 describes related work on legal domain ontologies concerning the privacy and protection of data. Section 4 outlines the functional and non-functional requirements. The methodology for developing the ontology is discussed in section 5. Section 6 outlines the design, implementation, and evaluation of a prototype POPIA ontology and the conclusion and future development follow in the last section.

2. Semantic Technology

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. In the context of the World Wide Web (WWW), search engines have made rapid progression in handling and providing vast amounts of information. However, search engines are struggling to identify relevant results for search terms (Benjamins et al., 2002). Semantic technology is regarded to be the single most important factor for advancing the Web; it encodes meanings separately from data. It is able to deal with large data sets and link them together via self-describing interrelations, allowing it to be processed by machines. It is viewed as the leading framework to deal with the diverse and huge size of assets on the Web. The Semantic Web is an expansion of the present Web, wherein data is given unambiguous meaning. (Sheth et al., 2004). The ability to make good decisions efficiently is essential in allowing organisations to better manage and process large amounts of data.

An ontology is a formal representation of the knowledge by a set of concepts within a domain and the relationships between those concepts (Man, 2013). It is the process of formalising 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 (Noy & McGuinness, 2001). Ontologies can represent vocabularies to investigate the connections between various terms, enabling machines and people to decipher their "meaning" unambiguously (Horrocks et al., 2003). 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*" (Horrocks et al., 2003). As a result, ontologies introduce a sharable and reusable knowledge base, enabling the extension of knowledge of a given domain.

The formal representation of the information in an ontology allows for the extension of existing knowledge via inferences made on the existing knowledge base. The reasoning capabilities associated with technologies also allow for the identification of data inconsistencies in the knowledge base. Ontologies are used in various domains as a type of information depiction about the world or some subset of it (Man, 2013). Examples of domains include Artificial Intelligence, Cybersecurity, the Semantic Web, Biomedical Informatics and the legal domain. Ontologies in the legal domain support the organisation of legal documents and provide support for legal reasoning (Bartolini & Mathuri, 2015a). Several successful ontologies have developed in the health sector domain, for example, the Gene Ontology (Consortium, 2001) and the Protein ontology (Sidhu et al., 2005).

The Web Ontology Language (OWL) is a "Semantic Web language intended to represent and capture rich and complex knowledge about things, and their relations between them" (Parsia et al., 2012). The knowledge expressed by OWL enables it to be reasoned about by using automated reasoners that verify the consistency of the domain knowledge within an ontology or revealing hidden knowledge. OWL ontologies promote reuse and modularity, as it can be published in the WWW and be referenced from other OWL ontologies. (Parsia et al., 2012). Knowledge concepts captured from data, in a given domain, are reasoned about in a rich hierarchical structure of concepts and their inter-relationships (Sidhu et al., 2004). These relationships help with the matching of concepts even if the data sources describing these concepts are not 100% uniform. Modelling 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 used to make inferences from this knowledge" (Krötzsch, 2012). OWL 2 is a World Wide Web Consortium (W3C) standard. SPARQL, standardised by W3C (SeaBorne et al., 2008), is the standard query language for ontologies.

3. Related Work

The POPIA is very closely aligned with the GDPR, the latter being more expressive. 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 (Da Veiga et al., 2018). There are various approaches related to expressing jurisdictional regulations such as the GDPR as an ontology (Pandit et al., 2018, Miles et al, 2007). The development of a POPIA-based ontology is inspired by these ontologies; no such ontology could be found in the literature at the time of writing. The objective of this paper is the development of an ontology to demonstrate concepts expressed within the POPIA and how it will influence the data subjects and responsible parties upon enactment. The goal is to create a prototype knowledge base for the POPIA concerning the lawful processing of personal information. The source for the descriptions of the POPIA concepts is derived from the Act's documentation (South African Government Gazette, 2013).

Pandit and Lewis (2017) 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) (Miles et al., 2007) as a source for the descriptions of the GDPR concepts, and the developers described it as "... a Semantic Web language for representing formally structured vocabularies". The terms are linked to the appropriate focus points 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.

Bartolini et al. (2015b) developed an ontology for the GDPR for the data protection requirements. It shows the data protection prerequisites with regards to the GDPR changes and introduces a methodology for incorporating it into a workflow 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. This was done to create an ontological representation of the obligations of data controllers, and the comparing privileges of information subjects.

The GDPRov project (Pandit & Lewis, 2017) is an ontology concerned with the management of compliance through recognising provenance information identified with assent and individual personal information required for consistency documentation. It is an OWL 2 linked open data ontology 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". The ontology uses SPARQL, an RDF query language, to query the provenance information described to find information relevant for compliance.

The PrOnto ontology (Palmirani et al., 2018) is a privacy ontology that captures 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 modelling rights and duties". The objective of PrOnto is to assist with legal thinking and verification of compliance. It achieves this by applying defeasible logic reasoning as opposed to solely information retrieval. (Palmirani et al., 2018). Defeasible reasoning is a rule-based method for efficient reasoning with inconsistent data (Kontopoulos et al., 2013).

4. Requirements

4.1. Functional Requirements

The requirements for the development of a POPIA ontology are defined by a set of competency questions the ontology should be able to answer. These questions will serve as the litmus test in the evaluation phase of the future mature ontology and helped to define the scope of the ontology (Noy & McGuinness, 2001). A few of these questions are:

- What is the responsibility of the various role-players?
- What is considered to be personal information?
- How does a role player who collaborates with an international organization deal with the management of compliance with the Act?

The set of competency questions is likely to change as the POPIA ontology matures and is being evaluated and maintained.

4.2. Non-functional Requirements

Protégé (Musen, 2015), an open-source ontology editor created at the Stanford University Center for Biomedical Informatics Research, was used to develop the ontology. 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. One of the design goals of Protégé is to be compatible and adaptable with other systems for knowledge representation and knowledge extraction. (Lambrix et al., 2003). This compatibility enables possible integration with other knowledge bases within the same domain.

Since the ontology is based on an Act, it will always be susceptible to change, due to changes or amendments in regulations. An initial list of non-functional requirements that will need to be satisfied is Adaptability, Reusability, Configurability, Testability, Maintainability, and Quality.

5. Methodology

There are various methodologies to develop ontologies including the Cyc methodology, the methodology of Uschold and King, the methodology of Gruninger and Fox, METHONTOLOGY, the KACTUS approach, and the SENSUS-based methodology (Fernández-López et al., 2002). METHONTOLOGY enables the development of ontologies at the knowledge level and follows an iterative approach utilising evolving prototypes. It supports prototyping and comprises of the following processes: Specification, Knowledge Acquisition, Conceptualisation, Formalisation, Implementation, and Maintenance. Noy & McGuinness' *Ontology Development Guide* (2001) presents an iterative development process that repeats continuously to enhance the ontology and 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.

The authors used a combination of METHONTOLOGY and the Noy & McGuinness methods. The two methods have complementary processes. METHONTOLOGY provides high-level activities for the development life cycle and the *Development Guide* method provides granular steps for design and implementation which are intricate phases in the development life cycle.

METHONDOLOGY was implemented in the higher-level development processes such as the specification and knowledge acquisition processes. During the specification phases the purpose of the ontology, intended users and scope were determined. The output of the specialisation phase is a natural language description. The knowledge acquisition phase consisted of finding sources of information such as the results on GDPR ontology development, and occurs concurrently with the first phase. The other phases in METHONTOLOGY overlap to a great extent with those of the Noy and McGuinness methodology but the latter is more specific. The next section explains how the Noy & McGuinness methodology was followed to design the POPIA ontology.

6. Design, Implementation & Evaluation

The design follows the sub-processes of the *Ontology Guide* as specified in section 5. All the key terms considered important for the knowledge base are listed, an initial class hierarchy was defined, and properties for classes and their features were identified.

6.1. Important Terms

The identification of terms from the Act will assist in explaining concepts in the POPIA. The terms were 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 (Noy & McGuinness, 2001). 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 conceptualise and establish the class definitions of the ontology.

6.2. Classes

The Conditions for lawful processing of personal information is defined in Chapter 3 of the Act. 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 the purpose and means of the processing and during the processing itself" (South African Government Gazette, 2013).

The list of conditions is "Accountability, Processing limitation, Purpose specification, Further processing limitation, Information quality, Openness, Security Safeguards and Data subject participation". A subset of these concepts was used to define the classes for the lawful processing of personal information. From this list, the following subclasses of the *LawfulCondition* root class were created: *Accountability, InformationQuality, Lawfulness, Openness, PurposeSpecification, and Security*.

The rights of data subjects is defined Chapter 2 Section 5 of the Act. "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". 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". (South African Government Gazette, 2013). From these rights, a *DataSubjectRight* class was created to with subclasses representing the individual rights: *RightToCorrection, RightToDelete, RightToObject* and *RightToSubmitComplaint*.

Data processing (Chapter 1) provides the definitions and purpose of the Act. 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" (South African Government Gazette, 2013). This includes amongst others the collection, storage, transfer or destruction of information. In addition, 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". This section highlights automated processing and profiling which are used to further describe data processing.

These concepts are used to extend the class definitions for conceptualising data processing: A *DataProcessing* class with subclasses *ProcessingActivity, AutomatedProcessing, InformationTransfer, Profiling, and ManualProcessing*.

Person is described in Chapter 1 of the Act as "a natural person or a juristic person". It also 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". 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". A data subject refers to "the person to whom the personal information relates". An operator is "a person who processes personal information for a responsible party". The "responsible party is a public or private body or any other person which, determines the purpose of and means for processing personal information".

Finally, the last concept identified for defining the “person” concept is the regulator. (South African Government Gazette, 2013). 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*.

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.

Table 1. Core class Hierarchy

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

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

- 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.

The highlighted terms are used to create properties. For example, “exercise rights” relates 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.

Table 2. Core properties

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

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. (2015b) were used, and in some instances changed, to fit into the POPIA ontology. The GDPR ontology 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 were used: *accessData*, *exerciseRight*, *objectTo*, *performProcessing* and *processingPerformedBy* (Bartolini et al., 2015b).

6.4. Evaluation

The ontology was evaluated by performing a set of queries using the DL (Description Logic) Query plugin in the Protégé editor. The evaluation consisted of answering an initial set of competency questions which includes querying the rights of data subjects, describing personal information, special personal information and the responsibility of the responsible party.

To provide an overview of the ontology model, in terms of numbers, the following statistics is provided:

```
Classes           : 60
Object properties : 13
Data properties   : 3
Individuals       : 7
Nodes             : 65
Edges            : 119
```

The high number of edges provides an idea of how interconnected the classes are. Higher numbers, generally allows for further inferencing. This allows more general queries to return possible individuals without having to define sub-relations explicitly.

Figures 1 to 5 show two queries. Both the individuals, *UWC* and *Teenager* are members of the *Person* class; *UWC* is a member of the subclasses *ResponsibleParty* and *DataSubject* while *Teenager* is in the subclass *DataSubject* and the sub-subclass *ChildDataSubject*. Figures 1 and 2 show the representation of these individuals in the ontology respectively.

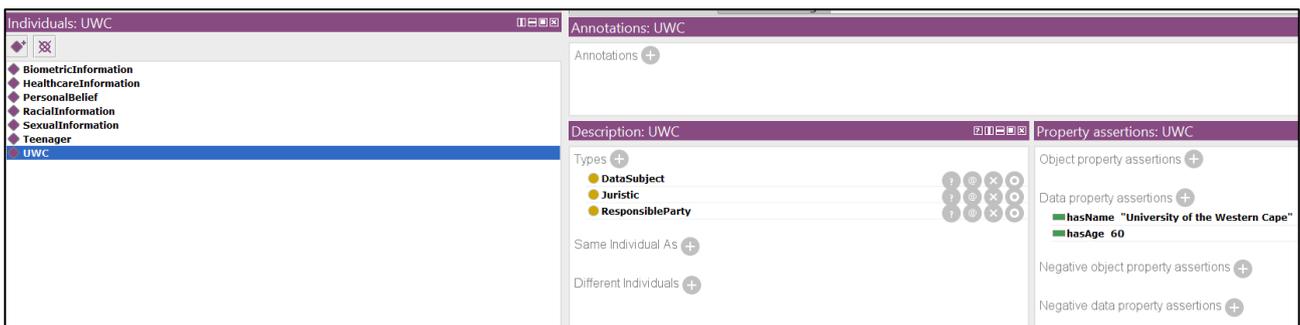


Figure 1: Description of the Individual "UWC"

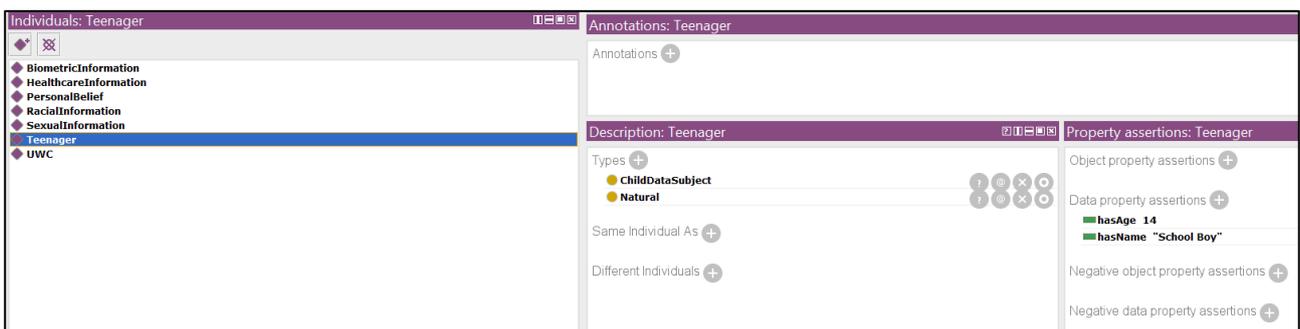


Figure 2: Description of the Individual "Teenager"

The first query is “List any entities who have to ensure that some condition required by the Act is satisfied AND has some PersonalInformation”. The query and the reply of the reasoner (i.e. *UWC*) is shown in Figure 3. Figure 4 shows the explanation given by the reasoner for the answer of the first query.

Then we change the query to “List any entities who have to ensure that some condition required by the Act is satisfied OR has some PersonalInformation”. The answer (i.e. *UWC* and *Teenager*) is shown in Figure 5.

DL query:

Query (class expression)

(hasInformation **some** PersonalInformation) **and** (mustEnsure **some** LawFullCondition)

Execute Add to ontology

Query results

Superclasses (4 of 4)

- DataSubject
- Person
- ResponsibleParty
- owl:Thing

Subclasses (1 of 1)

- owl:Nothing

Instances (1 of 1)

- ◆ UWC

Figure 3: The first query and answer

Explanation for UWC Type (hasInformation some PersonalInformation) and (mustEnsure some LawFullCondition)

Show regular justifications All justifications
 Show laconic justifications Limit justifications to

Explanation 1 Display laconic explanation

Explanation for: UWC Type (hasInformation some PersonalInformation) and (mustEnsure some LawFullCondition)

- UWC Type DataSubject
- UWC Type ResponsibleParty
- ResponsibleParty SubClassOf mustEnsure some LawFullCondition
- DataSubject SubClassOf hasInformation some PersonalInformation

Figure 4: Explanation of the answer to the first query

DL query:

Query (class expression)

(hasInformation **some** PersonalInformation) **or** (mustEnsure **some** LawFullCondition)

Execute Add to ontology

Query results

Superclasses (2 of 2)

- Person
- owl:Thing

Subclasses (6 of 6)

- ChildDataSubject
- Complain
- DataSubject
- Objection
- ResponsibleParty
- owl:Nothing

Instances (2 of 2)

- ◆ Teenager
- ◆ UWC

Figure 5: The second query and answer

8 Conclusion and Future Work

This paper highlights the importance of private information regulation in a globally connected world and 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 POPI Act and the impact it will have on various role players, such as the responsible party and data subjects. It outlines the development of a prototype 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 must be continued to produce a mature ontology. The intended users of this ontology (when it has matured) include any person or organisation that will be affected by the POPIA, for example, citizens, businesses, government departments, and community centres.

The POPIA ontology is based on similar work done for the European Union's General Data Protection Regulation (GDPR). This research demonstrates a proof of concept knowledge base on the POPIA. Going forward, this research can benefit by involving a legal domain expert in the evaluation and subsequent design decisions of the ontology. This will enable better accuracy in modelling decisions with regards to describing concepts formally. The reuse of existing relevant ontologies is a consideration for future work.

References

- Bartolini, C. & Mathuri, R. (2015a). Reconciling data protection rights and obligations : An ontology of the forthcoming EU regulation (2015). Workshop on Language and Semantic Technology for Legal Domain (LST4LD), Recent Advances in Natural Language Processing (RANLP 2015).
- Bartolini, C., Muthuri, R., Santos, C. (2015b). Using ontologies to model data protection requirements in workflows. Ninth International Workshop on Juris-informatics (JURISIN)
- Benjamins, V.R., Contreras, J., Corcho, O., Gomez-Perez, A. (2002). Six challenges for the semantic web. International Conference on Principles of Knowledge Representation and Reasoning.
- Birnhack, M.D. (2008). The EU Data Protection Directive: An Engine of Global Regime. *Computer Law & Security Review*, Vol. 26, no. 6. pp.508-520.
- Botha, J., Grobler, M.M., Hahn, J., Eloff, M.M. (2017). A High-level Comparison between the South African Protection of Personal Information Act and International Data Protection Laws, 12th International Conference on Cyber Warfare and Security (ICWS), Dayton, Ohio, USA.
- Common, M.F. (2018). Facebook and Cambridge Analytica: let this be the high-water mark for impunity (07 2018), <http://eprints.lse.ac.uk/88964/>
- Consortium, T.G.O. (2001). Creating the gene ontology resource: Design and implementation. *Genome Research* Vol. 11, no. 8, pp. 1425-1433. Retrieved from <https://doi.org/10.1101/gr.180801>, <http://dx.doi.org/10.1101/gr.180801>
- Constitution of the Republic of South Africa: Section 14, (1996). <http://www.justice.gov.za/legislation/constitution/SACConstitution-web-eng.pdf>
- Da Veiga, A., Vorster, R., Li, F., Clarke, N., Furnell, S. (2018). A comparison of compliance with data privacy requirements in two countries. 26th European Conference on Information Systems. University of Portsmouth.
- Fernández-López, M., Gómez-Pérez, A. (2002). Overview and analysis of methodologies for building ontologies. *Knowledge Engineering Review*. Vol 17, no. 2, pp. 129-156 (Jun 2002). Retrieved from <https://doi.org/10.1017/S0269888902000462>, <https://doi.org/10.1017/S0269888902000462>
- Horrocks, I., Patel-Schneider, P., Harmelen, F. (2003). From SHIQ and RDF to OWL: The making of a web ontology language. *SSRN Electronic Journal*, Vol 1. Retrieved from <https://doi.org/10.2139/ssrn.3199003>
- Information Shield. (nd). International privacy laws. <http://www.informationshield.com/intprivacylaws.html> [Accessed July/7, 2014]
- Kontopoulos, E., Bassiliades, N., Governatori, G., Antoniou, G. (2013). A modal defeasible reasoner of deontic logic for the semantic web. *Semantic Web*, pp. 140-167. <https://doi.org/10.4018/978-1-4666-3610-1.ch007>
- Krötzsch, M. (2012). OWL 2 Profiles: An Introduction to Lightweight Ontology Languages, pp. 112-183. Springer Berlin Heidelberg, Berlin, Heidelberg (2012). Retrieved from <https://doi.org/10.1007/978-3-642-33158-94>;
- Lambrix, P., Habbouche, M., Perez, M. (2003). Evaluation of ontology development tools for bioinformatics. *Bioinformatics*, Vol. 1, no. 12 , pp 1564-71 (2003)
- Man, D. (2013). Ontologies in computer science. *Didactica Mathematica*, Vol. 31, pp. 43-46.
- Miles, A., Pérez-Agüera, J.R.: Skos (2007). Simple knowledge organisation for the web. *Cataloging & Classification Quarterly*, Vol. 43(3-4), pp. 69-83. https://doi.org/10.1300/J104v43n03_04, https://doi.org/10.1300/J104v43n03_04
- Musen, M.A. (2015). The Protégé project: a look back and a look forward. *AI Matters* Vol. 1, no. 4, pp. 4-12. <https://doi.org/10.1145/2757001.2757003>

- Noy, F., Mcguinness, D. (2001). *Ontology development 101: A guide to creating your first ontology. Knowledge Systems Laboratory 32, Vol 01.*
- Palmirani, M., Martoni, M., Rossi, A., Bartolini, C., Robaldo, L.(2018). PrOnto: Privacy Ontology for Legal Reasoning: 7th International Conference, EGOVIS 2018, Regensburg, Germany, pp. 139-152. https://doi.org/10.1007/978-3-319-98349-3_11
- Pandit, H.J., Fatema, K., O'Sullivan, D., Lewis, D. (2018). GDPPrEXT - GDPR as a Linked Data Resource. *Lecture Notes in Computer Science*. In book: *The Semantic Web*. DOI: 10.1007/978-3-319-93417-4_31
- Pandit, H.J., Lewis, D. (2017). Modelling provenance for GDPR compliance using linked open data vocabularies. In: *PrivOn@ISWC*.
- Parsia, B., Patel-Schneider, P., Krötzsch, M., Rudolph, S., Hitzler, P. (2012). *OWL 2 web ontology language primer (second edition)*. Tech. rep., W3C (Dec 2012), <http://www.w3.org/TR/2012/REC-owl2-primer-20121211/>
- Seaborne, A., Prud'hommeaux, E. (2008). *SPARQL query language for RDF*. W3C recommendation, W3C <http://www.w3.org/TR/2008/REC-rdf-sparql-query-20080115/>
- South African Government Gazette: *Protection of personal information act (2013)*, Retrieved from <http://www.justice.gov.za/legislation/acts/2013-004.pdf>
- South African Government Gazette (2108). <https://opengazettes.org.za/gazettes/ZA/2018.html>
- Sheth, A., Ramakrishnan, C. (2004). *Semantic (web) technology in action: Ontology driven information systems for search, integration and analysis. IEEE Data Engineering Bulletin*, Vol. 26.
- Sidhu, A., S. Dillon, T., Chang, E., S. Sidhu, B. (2005) *Protein ontology development using OWL. OWLED*05 Workshop on OWL: Experiences and Directions*, Galway, Ireland.
- Srncova, Z., Babela, R., Mamrilla, R., & Balazova, Z. (2019). *GDPR implementation in public health. Int Health Journal*. Feb 2019 Issue. <http://ihjonline.com/issue/2019-02/gdpr-implementation-in-public-health>
- United Kingdom Government Gazette (1998). *Data Protection Act*. <http://www.legislation.gov.uk/ukpga/1998/29/contents>