NIS2Onto: an Ontological Representation of the NIS 2 Directive

Gianpietro Castiglione [a,*], Daniele Francesco Santamaria [a], Giampaolo Bella [a] and Gaetano Puccia [b]
[a] *Università di Catania, Catania, Italy*
E-mails: gianpietro.castiglione@phd.unict.it, daniele.santamaria@unict.it, giampaolo.bella@unict.it
[b] *Intrapresa S.r.l., Via Emanuela Loi, 15 - Ragusa, Italy*
E-mail: gaetano.puccia@intrapresa-it.it

Abstract. This paper presents NIS2Onto, an OWL ontology designed to model and manage the complexities of the NIS 2 Directive, aimed at bolstering cybersecurity across essential sectors in the European Union. NIS2Onto offers the ontology that translates the Directive's legal and technical requirements into an ontological format, facilitating improved compliance management and enhanced understanding among cybersecurity professionals, legal experts, and organisational stakeholders. Through the ontological representation of the NIS 2 entities, relationships, and obligations, NIS2Onto enables automated compliance verification, streamlined risk assessments, and effective policy implementation. Our evaluation employs both metrical and qualitative analysis through a real case study in order to witness the robustness and practical applicability of NIS2Onto. The ontology not only supports the accurate interpretation of complex legal texts but also aids in the systematic enforcement of cybersecurity measures. Furthermore, NIS2Onto's extensibility allows for integration with other regulatory frameworks, fostering a comprehensive and unified approach to cybersecurity governance.

Keywords: Cybersecurity, Ontology, NIS 2, Compliance

1. Introduction

Technological developments have become complex and intricate throughout time, leading to the settlement of institutions whose duty is to establish standards that regulate these advancements and design best practices for their adoption. Recently, this also happened in the area of artificial intelligence with the EU Artificial Intelligence Act (AI Act), which inevitably involved the domain of security and privacy. In contrast to concrete standards, which recommend using a particular technology over another and take a systems-oriented approach, the various directives and regulations that are promulgated in the security and privacy realms end up being even more important because the primary objective is to safeguard human activity.

Among the most recent directives promulgated by the European Union, the NIS 2 Directive essentially establishes measures that operators of technical processes and national states must abide by. The NIS 2 Directive is an advancement over the NIS Directive, which was issued less recently. It covers a wide range of topics, from the one security measure that should be implemented on an IT system to the coordination requirements that member states must follow, and the different guidelines that dictate what should be done in the event of a data breach and the notifications that follow.

Hence, starting from the NIS 2 Directive, the result of the present contribution is the OWL ontology called NI2Onto, a complete ontological representation of the NIS 2 Directive, reachable at [10]. Representing a Directive

[*] Corresponding author. E-mail: gianpietro.castiglione@phd.unict.it.

through an ontology means extrapolating the security measures, deconstructing the agents, actions, and objects of such measures, and associating them with the various and appropriate ontological commitments. The objective of NIS2Onto is to provide a semantic representation of a complex security document, now searchable even only in correspondence with certain agents and measures. Our ontological representation explains the structure of the document and enhances it with automatic reasoning capabilities, a decisive outcome for the compliance verification process, particularly of large enterprises and organisations.

In short, NIS2Onto reduces time and effort required for the verification, both by enabling continuous monitoring of any modification that occurs during the lifespan of the company and by minimising the risk of error by human personnel. This also implies that fewer resources are required for compliance verification, leading to cost savings and increasing efficiency. Moreover, a wide range of compliance requirements can be covered, ensuring that no aspect is overlooked: these requirements can change over time and such changes are handled automatically without human effort. Finally, automatic compliance verification makes the audit process fast, easy, and detailed: this is crucial for companies and organisations that must demonstrate to regulators that they adhere to security measures.

Coherently with what is outlined above, our work is supported by an IT enterprise, Intrapresa S.r.l., [1] operating in the field of innovative software for petrol stations. The real-world case study to which NIS2Onto is applied below comes from this application domain. Worthy of mention, the present work extends a previous contribution [2] by the same authors, where the ontological approach for the NIS 2 Directive is just sketched out.

The paper is organised in the following sections: related work in Section 2; NIS2Onto ontology in Section 3 together with a metric evaluation; a case study illustrating the value, concreteness, and applicability of the ontology is shown in Section 4; the paper concludes in Section 5 with some final remarks and future prospective.

2. Related work

Ontologies are largely employed in a wide range of scientific domains and have been developed substantially. The same stands for cybersecurity — with a focus on legislation — thus we concentrate on the main contributions that are spendable in such realm.

Fenz [8] aims to introduce a method for formalising information security control descriptions and a decision support system that enhances automation, thereby improving the cost efficiency of the information security compliance checking process. The author applied this method to ISO 27002 information security controls and created a semantic decision support system.

Tailhardat et al. [19] developed NORIA-O ontology, to define an infrastructure, its events, diagnostics, and remediation operations carried out during incident management. An example use case detailing a hypothetical failure illustrates how this ontology may be used to describe intricate scenarios and provide a foundation for anomaly identification and root cause investigation.

Syed et al. [18] presented the UCO ontology. UCO is designed to help cybersecurity systems with information integration and cyber situational awareness. To facilitate information sharing and exchange, the ontology combines and integrates diverse data and knowledge schemas from many cybersecurity systems with the most widely used cybersecurity standards.

Muñoz et al. [14] developed the PPROC ontology, which is designed to provide semantic descriptions of public procurement contracts and procedures, therefore supporting both disclosure and accountability. The PPROC ontology is comprehensive as it includes information on the whole process, from the first contract publication to its termination, in addition to the standard information on the tender, its goals, deadlines, and recipients.

Syed [17] presented a conceptual model for formal knowledge representation of the vulnerability management area, which is the Cybersecurity Vulnerability Ontology (CVO). CVO is used to create a Cyber Intelligence Alert (CIA) system that sends out cyber warnings on vulnerabilities and countermeasures.

Wang et al. [20] developed OVM, which is composed of the vulnerabilities in NVD, for vulnerability management (OVM), along with extra inference rules, knowledge representation, and data-mining techniques. OVM offers a

[1] https://www.intrapresa-it.it/

potential road to the success of ISAP through the smooth integration of common vulnerabilities and their associated concepts, such as attacks and solutions.

In order to provide a legal knowledge modelling of the privacy agents, data categories, kinds of processing activities, rights, and duties, Palmirani et al. [15] developed PrOnto, a legal ontology on the GDPR. This approach is based on ontological patterns combined with legal theory analysis.

In the work of Elluri et al. [6] the regulations required by PCI DSS and EU GDPR have been represented by an integrated, semantically rich knowledge graph. The authors have recognised the responsibilities outlined in these rules and connected them to relevant Cloud Security Alliance (CSA) controls in the Ontology.

Always with the idea of managing the difficulties of complex documents, Joshi et al. [11] created a semantically rich ontology and created a database with many policy documents as instances of this ontology. Using deontic logic, the authors identified rules from these policy documents that may be utilised to automate data privacy management.

Ontologies have been applied to specify anonymisation policies. A general anonymisation policy for usage in big data and cloud platforms is presented in the study of Matsunaga et al. [12]. A proposed ontology with precise formal requirements for data anonymisation procedures is also presented to standardise the application of data anonymisation policies and facilitate the reuse of data anonymisation policies.

As far as we know, representing cybersecurity documents with the goal of compliance verification, particularly for the NIS 2 Directive, is an uncharted path. Therefore, the effort of the present contribution: exploiting a general yet practical and concrete approach to represent and operationalise security directives, ultimately enhancing the effectiveness and efficiency of the compliance process.

3. The NIS2Onto Ontology

3.1. Overview of the NIS 2 Directive

The NIS 2, Directive (EU) 2022/2555 [7], was adopted on 14 December 2022, and is effective from 16 January 2023, replacing Directive (EU) 2016/1148. The Directive sets rules for security risk management across the nations of the European Union regarding the most important sectors. In some cases, it aligns with related legislation and defines how nations have to cooperate through the establishment of specific groups of relevant stakeholders. The member states, namely, the members of the European Union to which the NIS 2 is applied, will have to adopt the new version of the Directive within 21 months. The document is constituted of seven chapters, organised in articles and then in paragraphs. Chapters I and VIII contain general considerations on the applicability of the Directive and are not relevant to our purpose.

The measures are both organisational and security-related; the former defines how the various entities should act, including peers to cooperate with and timelines; the security measures, which are actually a small part, are conceived for *important* and *essential* entities and define some security best practices. Important entities are those companies operating in critical sectors or of medium-size (50–249 employees or a turnover between ten and fifty million euros and 50M) operating in any of the sectors in the scope of the NIS 2. On the contrary, essential entities are large companies operating in sectors of critical importance (more than 250 employees or more than fifty million euros in turnover).

3.2. SecOnto Methodology

NIS2Onto is a comprehensive ontological representation of the European NIS 2 Directive by fully leveraging the SecOnto methodology[5]. Specifically, SecOnto re-examines the phases of Methontology [9] in light of domain-specific factors. SecOnto anticipates five phases. *Preprocessing* involves a broad evaluation of the document and its content in order to assess and individualise the structure of the legal document as well as the many ways that security measures are organised. One of the results of the step is a broad comprehension of the document. It somewhat overlaps with step 2 of knowledge acquisition and predicts the first step of Methontology, the *Specification* phase, since a thorough examination of the legal documents is necessary to completely comprehend the field of knowledge and gather all the necessary data. SecOnto then recommends the *Interpretation* stage to determine the elements of the measurements, such as the subjects, acts, and objects (i.e., the ontological triples). These are the principal results of the step. It combines the first two processes of Methontology, namely, *Specification* and *Knowledge Acquisition*.

Since the *Structuring* stage of SecOnto offers the initial semi-informal definition of the ontological model and the necessary knowledge — which are also the key outcomes — it combines the steps *Conceptualisation* and *Integration* of Methontology. SecOnto then offers the *Representation* stage, which is the complete ontological model translation of security procedures. This stage is associated with the *Implementation* step of Methontology. The main results are the ontology that describes the measurements, the associated paperwork, and the conclusions drawn from the ontology. SecOnto ends with the *Verification* phase, which shows how to use the structures that were produced for compliance verification using ontological reasoning and how to use SPARQL language for differential analysis of post-compliance verification. This step combines the *Evaluation* step of Methontology with the measures' compliance verification. The ontology verification, the SPARQL queries confirming the measures' compliance, and the compliance verification itself are the results of such steps. Every step consists of many micro-steps, each one endowed with a description, an italicised brief introduction, and a real-world example to illustrate how to perform the micro-step.

Hence, by applying the *Preprocessing* step of SecOnto, we identify the chapters from II to VII, composed of the articles from 7 to 37, as the most suitable for the ontological representation. Since NIS2Onto is specific to the domain of security directives of the NIS 2 and tailored to the sections where the security measures are expressed, no foundational ontology has been initially considered for the development of NIS2Onto, even though some are undoubtedly useful for the integration of the work with other contributions. At the current stage of development, these foundational ontologies provide no significant contribution to either the methodology or the concrete application. Additionally, most of them are mainly devised for contexts outside cybersecurity, while, to the best of our knowledge, there are no significant contributions aimed at representing security documents.

Some ontological efforts rely on technological systems and processes, but these cannot be taken into account for the current stage of the development of NIS2Onto. This does not exclude that upcoming versions of NIS2Onto may adapt to domain-related ontologies, even in correspondence with the evolution of the NIS 2 Directive itself.

3.3. Automating the Creation of the Ontology

Creating the ontology was a challenging undertaking, mainly due to the complexity of the legal terminology and the complex connections between the elements in the security measures. Given these challenges, our initial strategy involved the application of automation techniques, with a particular focus on leveraging the capabilities of various natural language processing (NLP) tools for the extraction and accurate tagging of parts of speech (POS). This automated approach partially allowed us to efficiently process the complex legal language and lay a solid foundation for the subsequent steps of ontology development. The approach, which we previously proposed [4], saw the application of SpaCy and ClausIE library for the extraction of the relevant part of a document sentence and then of a security measure. By evaluating the resulting percentages [4], we can conclude that a complete and correct extraction of the elements did not take place. Considering the prospect of a fully automated ontological representation depending only on the extracted POS would have presented significant limitations as 100% accuracy would have been required in all cases. However, the ontology is generated semi-manually, since the correctly extracted POS were used for the initial creation of entities and properties after manual evaluation.

We deliberately avoided using certain tools, particularly in the context of ontological automated development, because they proved to be ineffective or incompatible with the context of security directives.

3.4. Competency Questions

NIS2Onto [10] is built to answer the following type of competency questions, which are translated in the corresponding SPARQL queries as shown in Section 4.

1. **Compliance check**. We want to check for the compliance of a certain individual, or an individual of a certain type, or the compliance of an individual with respect to a specific article.
2. **Differential analysis I**. We want to verify or search for the security measures that involve some entity (undertaking vs. entity).

3. **Differential analysis II**. We want to verify or search for the agents that are involved in some security measures (undertaking vs. important or essential).
4. **Specific search I**. Given an object (e.g., the risk assessment), we want to search for the article that threatens it.
5. **Specific search II**. Given an object (e.g., the risk assessment), we want to search for the action that involves it.
6. **Specific search III**. We want to search for all objects that syntactically contain a certain word in order to verify which elements are involved.
7. **Specific search IV**. We want to search for the standard associated with a specific category (e.g., vulnerability assessment, risk assessment, inconsistencies management).
8. **Integration I**. We want to integrate the measures of common agents among other security directives/regulations, e.g., GDPR.
9. **Integration and Differential Analysis**. We want to check the security measures missed by an agent across several directives/regulations.

Below we briefly present the constitutional elements of NIS2Onto and then introduce the evaluation of the ontology.

3.5. NIS2Onto Overview

Based on the SecOnto methodology, NIS2Onto inherits its characteristics as follows. In particular, NIS2Onto takes the security measures described in the NIS 2 Directive and encodes them ontologically according to the OWL specifications. Following a purely grammatical and then subsequently semantic structuring, in NIS2Onto, in each security measure, there are agents responsible for the security measure (who must carry out the action described in the security measure identified by the subject), the action to be carried out to comply with the security measure (the verb of the security measure), and the object or series of objects of the security measure (to whom the security measure is directed, in terms of either further agents or further processes/elements in general, identified by the object of the security measure).

NIS2Onto reflects the principle of SecOnto that associates specific agents with the security measures the agents must fulfil through the equivalence relationship, i.e., *EquivalentTo*, with a class description devising the features of the security measures, the object properties representing the actions and entities representing the object of the actions. Since NIS2Onto reflects the document-oriented design from SecOnto, this step is done for each security measure. From this representation, we can refer to both the security measure of the document, together with the related articles, chapters and paragraphs, and to the agent and the related actions and recipients. This is particularly useful to keep the ontology aligned with the document to identify the original location of the measure and its description and to enable a semantic search of its constitutional elements.

The provided representation of the security measures offers a dual standpoint for the critical and crucial task of compliance verification. This is possible thanks to the thoroughness of the domain, allowing for the use of both a top-down and bottom-up approach in evaluating individual instances. The first approach is based on the created ontology without any modification of its classes or properties and essentially follows exactly what is shown in the methodology. It consists of creating the individuals to be verified, assigning the security measures to them, again using the instantiating of classes and properties, and then evaluating the inferences generated by the reasoning. Conversely, the second approach takes a divergent path by leveraging and expanding upon the NIS 2 notions of *sector* and *entity*, which hold pivotal significance in achieving the objectives outlined in NIS 2. Notably, NIS 2 is explicitly geared towards addressing entities considered as crucial components within a nation's critical sectors. This approach delves deeper into ontology development by introducing new classes encapsulating the most pertinent existing standards. This distinctive approach is tailored to foster a context-specific evaluation of the business landscape, with a primary focus on the standards currently in practice. It is essential to note that the NIS 2 Directive refrains from imposing the adoption of specific standards, providing each enterprise with the autonomy to make its own choices. Consequently, referencing the measures of a directive can initially appear abstract in the absence of classes that represent concrete standards. The fundamental premise here is that when a specific standard is adopted, adherence to the security measures associated with that standard is inherently guaranteed, thus substantiating compliance with those measures. Building upon this concept, the idea materialises in creating these missing classes. For instance, one might have a subclass like *PCI-DSS* for class *Multi-factor Authentication* or subclass *NIST SP 800-53* for class *CryptographyPolicy*. With the introduction of such classes, the compliance verification task is twofold. Initially, it involves a more intricate internal audit where the specific standards in use are systematically categorised. Subsequently, harnessing the power of inferences, these specialised classes become directly associated with the broader concept expressed within a security measure. This approach, which aligns compliance with specific, adopted standards, further refines the verification step and empowers a more characteristic understanding of adherence to NIS 2 measures.

3.6. Classes and individuals

As stated before, the goal of NI2Onto is to represent the agents and objects of the security measures, keeping intact the structure of the document. For this purpose in NI2Onto, there are four main conceptual families of classes:

Document These classes are useful to coherently maintain the document structure and associate the articles and the paragraphs, eventually, to the related agent. Chapters are organised in articles, articles in paragraphs, where each paragraph refers to specific agents, actions and objects. The structural organisation of the document is reflected in a suitable hierarchy of classes, thus allowing the association of agents to the security measures in the document where they are mentioned.

Agent For the purpose of the NIS 2 Directive, it is crucial to identify the agents involved in the provided security measures. NIS2Onto has been developed with a focus on this goal, in particular, on the agents relevant to the NIS 2 Directive, for which the compliance should be verified. Therefore, the ontology provides suitable classes, one for each type of agent, where the related instances model the agents depicted in the Directive.

Object As in the case of the agents, the objects of the security measures are introduced by way of a suitable hierarchy. In most cases, it is hard to identify the object of a security measure and to associate it with a specific category. With the support of automatic tools, the classes representing objects have been generated from the document. The class names adopted has been obtained by removing from the object name the articles, adverbs, and other grammatical structures.

Compliance These classes are introduced to infer to which article an entity is compliant. Each class is suitably defined to check for the measures that must be satisfied by the compliant entity. For example, let us consider an excerpt of Article 10 from the NIS 2:

> **Computer security incident response teams (CSIRTs)**
>
> 1. Each Member State shall designate or establish one or more CSIRTs. The CSIRTs may be designated or established within a competent authority. The CSIRTs shall comply with the requirements set out in Article 11(1), shall cover at least the sectors, subsectors and types of entity referred to in Annexes I and II, and shall be responsible for incident handling following a well-defined process.
>
> 2. Member States shall ensure that each CSIRT has adequate resources to carry out effectively its tasks as set out in Article 11(3).
>
> 3. Member States shall ensure that each CSIRT has at its disposal an appropriate, secure, and resilient communication and information infrastructure through which to exchange information with essential and important entities and other relevant stakeholders. To that end, Member States shall ensure that each CSIRT contributes to the deployment of secure information-sharing tools
>
> ...

To verify the compliance of a member state with Article 10, the member state should satisfy each of its paragraphs. On these premises, it can be defined as an instance of an *Article-10-MemberState-Compliant*. We define such class as a subclass of:

```
Art10Par1-MemberState and Art10Par10 and Art10Par2
 and Art10Par3  and Art10Par4 and Art10Par5 and Art10Par6
 and Art10Par7-MemberState and Art10Par8 and Art10Par9
```

If we inspect the individual representing paragraph 3, namely *Art10Par3*, we notice that its compliance derives from satisfying the following conditions:

```
MemberState
 and ((allowTo some EntityAndStakeholderInformationExchange)
 and (ensureCSIRTContributesTo some InformationSharingToolsDeployment)
 and (ensureCSIRTHas some CommunicationInformationInfrastructure))
```

Analogous definitions apply to the other paragraphs and to the other articles of the NIS.

3.7. Object-properties

In NIS2Onto, the object properties are used to identify the actions taken by an agent with respect to some object. In the context of a security measure, the object property is introduced by the main verb of the measure, namely, the verb strictly related to the agent of the selected measure. Since there are different types of actions, we adopt the following name convention.

- The object property name reports the verb itself. This case occurs when the object is well-established, and it is not necessary to mediate by other terms.
- The object property name is obtained by the verb, but it includes other additional elements. This occurs when the action is complex or the verb is phrasal. For example, *allow* falls in the first case, while *allowTo*, *allowToAccess*, *allowToEntitiesTimeTo* fall in the second one.

3.8. Data-properties

In the NIS 2 Directive, numerical values can be expressed as objects of some actions, e.g., for establishing data breach notification months, or the period for submitting the national cybersecurity strategy. In such cases, data properties are adopted to express the actions with the same name convention adopted for the object properties.

3.9. SWRL rules

SWRL rules are introduced in NIS2Onto to extend the reasoning capabilities of the ontology and to enforce compliance verification. We identify the following two scenarios:

Self-referencing entities The first scenario regards the presence of self-referencing entities inside a security measure. For instance,

$$MemberState(?x) \wedge CompetentAuthority(?y) \wedge designate(?x;?y) \wedge \\ PointOfContact(?z) \rightarrow isSinglePointOfContact(?z,?x)$$

ensures that a security measure must be linked with the specific *Member State* to which it pertains. This approach becomes valuable during the compliance verification step. For instance, if multiple *Member States* are instantiated and associated within the related security measure, the rules will trigger the inconsistency. Essentially, this rule reinforces the connection between the measure and the specified *Member State*.

Describing preconditions and implications The second scenario concerns specific measures of the Directive that might encompass structures that involve not only explicit security measures but also conditional elements that prescribe the application of these measures. These conditional clauses introduce intricacies into both the interpretation of the Directive and the representation of such measures. As an example, we consider the following measure (Article 26, third paragraph, first sentence):

> "If an entity as referred to in paragraph 1, point (b), is not established in the Union but offers services within the Union, it shall designate a representative in the Union."

The measure is guaranteed by the following rule:

$$Entity(?\,x) \wedge NotEuropeanUnion(?\,x) \wedge referredTo(?\,x, ?\,y) \wedge offerServicesTo(?\,t, ?\,x) \wedge$$
$$UnionRepresentative(?\,u) \wedge EuropeanUnion(?\,t) \wedge Art23Par3 - b(?\,y) \rightarrow designate(?\,u, ?\,x)$$

We can now briefly discuss the evaluation of the ontology.

3.10. Evaluation

Metrics serve as guidelines for evaluating the complexity, quality, and usability of ontologies, facilitating their development and refinement, additionally offering insights into characteristics such as the number of classes, properties, instances, and logical axioms. Ontological metrics enable practitioners to gauge the structural richness and coherence of an ontology. This not only aids in the comparison and selection of ontologies for specific applications but also supports the iterative improvement of ontological models, ensuring they effectively represent the intended domain knowledge.

On the current version of the ontology, the following metrics derived from OntoMetrics [13] have been computed:

Base Metrics
- Axioms: 3574
- Logical axioms count: 1740
- Class count: 1330
- Total classes count: 1330
- Object property count: 363
- Total object properties count: 363
- Data property count: 24
- Total data properties count: 24
- Properties count: 387
- Individual count: 44
- Total individuals count: 44
- DL expressivity: ALCHQ(D)

Class Axioms
- SubClassOf axioms count: 1280
- Equivalent classes axioms count: 359
- Disjoint classes axioms count: 2
- GCICount: 0
- HiddenGCICount: 356

Data Properties Axioms
- SubDataPropertyOf axioms count: 10
- Data property domain axioms count: 1

Individual Axioms
- Class assertion axioms count: 43
- Object property assertion axioms count: 45

Annotation Axioms
- Annotation assertion axioms count: 72

Schema Metrics
- Attribute richness: 0.018045
- Inheritance richness: 0.962406
- Relationship richness: 0.361277
- Equivalence ratio: 0.269925
- Axiom/class ratio: 2.687218
- Class/relation ratio: 0.663673

Knowledge-base Metrics
- Average population: 0.033083
- Class richness: 0.030827

The evaluation of the ontology focuses on qualitative rather than quantitative aspects since the metrics associated with the ontology could be quite variable in correspondence with its possible modifications and changes. Although such assumptions may not seem entirely accurate in an ontological context, their justification derives from the nature of NIS2Onto, namely, a domain-related ontology that may not be fixed over time. For example, if a company wanted to use NIS2Onto to verify compliance with its requirements, and wanted to break down an object of a security measure into smaller objects, then it would be able to do so, without however destroying the meaning of the structures provided. It is therefore clear that the metrics would no longer be consistent if such a scenario occurred. Instead, the evaluation of the current version of NIS2Onto follows some best-known approaches classified by Raad et al. [16]:

- **Corpus-based**: *"Corpus-based approaches are used to evaluate how far an ontology sufficiently covers a given domain"*. NIS2Onto is an application-based ontology. It represents a specific context as it is created ad hoc on the NIS 2, so it entirely covers the domain of the Directive.
- **Task-based**: *"Task-based approaches try to measure how far an ontology helps to improve the results of a certain task"*. The principal aim of this paper is to provide a methodology for a representation that can be easily intelligible, make the entities and properties reusable with the integration of further ontologies, and provide mathematical support for compliance verification. Although there may be tools that assist with compliance

verification, they may not accomplish the peculiarities on which the ontologies are based, and which have already been contextualised.

- **Criteria-based**: *"Criteria-based approaches measures how far an ontology or taxonomy adheres to certain desirable criteria"*. The methodology is particularly suitable for obtaining structures that provide a representation that is accurate, concise, complete, efficient, and more generally respectful of the FAIR principles.

4. Case study

This section shows how to apply NIS2Onto to a real case study developed in the context of the business activities of Intrapresa S.r.l. In the considered case, the company would check its compliance with the NIS 2 Directive. Of course, this Section does not show the actual company's compliance obligations to preserve its security and privacy — the data that is shown is appropriately anonymised, namely, randomised by permuting the values of the obligations.

By following the ontology specification, we create a fresh individual that represents the company, and we use it as an interactive auditing tool. By exploiting the inferences obtained from the reasoning task on the ontology, we verify which articles the company is compliant to. As a first step, we first introduce the company as an instance of *Entity*, and then we describe all the features representing the company, namely, all the obligations the company currently satisfies. These are depicted in Figure 1. For example, since the company applies some strategies of disaster recovery, we state that it "include DisasterRecovery".

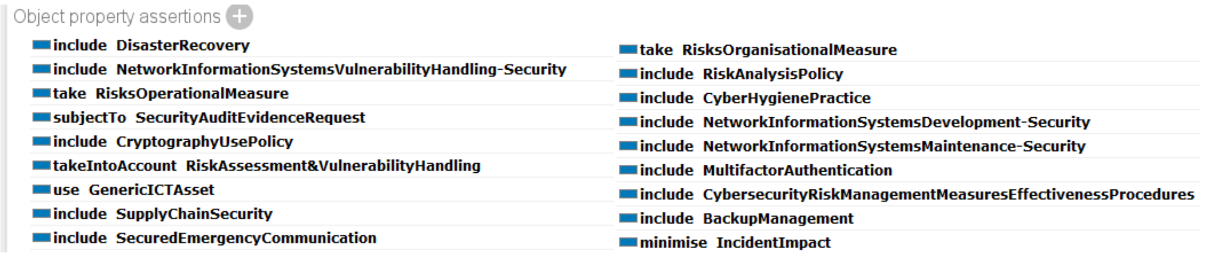

Object property assertions

include DisasterRecovery
include NetworkInformationSystemsVulnerabilityHandling-Security
take RisksOperationalMeasure
subjectTo SecurityAuditEvidenceRequest
include CryptographyUsePolicy
takeIntoAccount RiskAssessment&VulnerabilityHandling
use GenericICTAsset
include SupplyChainSecurity
include SecuredEmergencyCommunication
take RisksOrganisationalMeasure
include RiskAnalysisPolicy
include CyberHygienePractice
include NetworkInformationSystemsDevelopment-Security
include NetworkInformationSystemsMaintenance-Security
include MultifactorAuthentication
include CybersecurityRiskManagementMeasuresEffectivenessProcedures
include BackupManagement
minimise IncidentImpact

Fig. 1. Set of obligations satisfied by the company.

Then, we exploit the reasoning capabilities of the ontology in order to find out which articles the company is compliant with, and we can leverage them to understand which obligations are not satisfied, if any. In our case study, the company is currently satisfying the compliance measures partially shown in Figure 2.

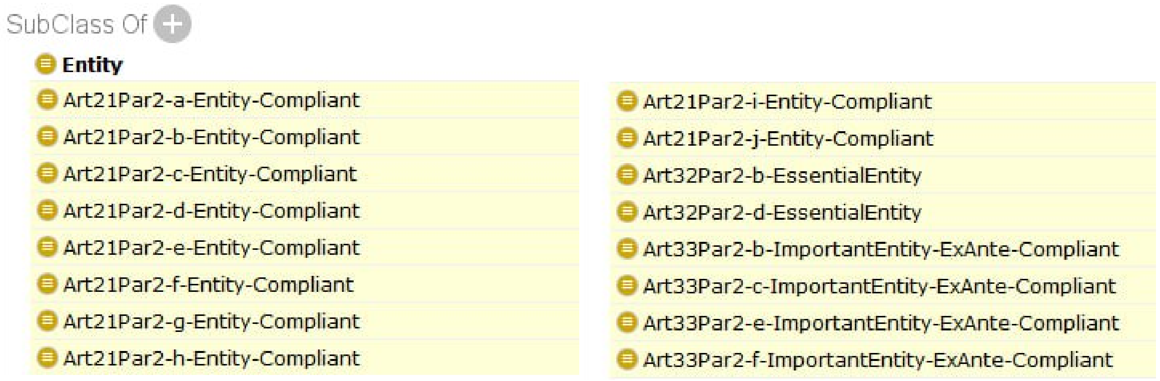

SubClass Of

Entity
Art21Par2-a-Entity-Compliant
Art21Par2-b-Entity-Compliant
Art21Par2-c-Entity-Compliant
Art21Par2-d-Entity-Compliant
Art21Par2-e-Entity-Compliant
Art21Par2-f-Entity-Compliant
Art21Par2-g-Entity-Compliant
Art21Par2-h-Entity-Compliant
Art21Par2-i-Entity-Compliant
Art21Par2-j-Entity-Compliant
Art32Par2-b-EssentialEntity
Art32Par2-d-EssentialEntity
Art33Par2-b-ImportantEntity-ExAnte-Compliant
Art33Par2-c-ImportantEntity-ExAnte-Compliant
Art33Par2-e-ImportantEntity-ExAnte-Compliant
Art33Par2-f-ImportantEntity-ExAnte-Compliant

Fig. 2. Compliance measures satisfied by the company.

When the company satisfies specific measures, the corresponding individual can be inferred as the NIS 2 agent prescribed to respect those measures. In our case study, the reasoner infers that the company is an instance of

ImportantEntity-ExAnte, since the individual is compliant with all the measures prescribed to a NIS 2 Important Entity; additionally, it is also *Ex-ante*, a term used to introduce those entities whose security measures are defined before a security incident occurs.

In case there are some missing features, we can exploit those measures to carry out a differential analysis to understand which articles the company is not compliant with. In our case, we perform the query corresponding to competency question 2 of Section 3 to discover which measures the company is not compliant with.[2] We can also discover which actions should be taken by the company in order to fulfil such measures and therefore be fully compliant with the NIS 2 Directive. The results are shown in Figure 3.

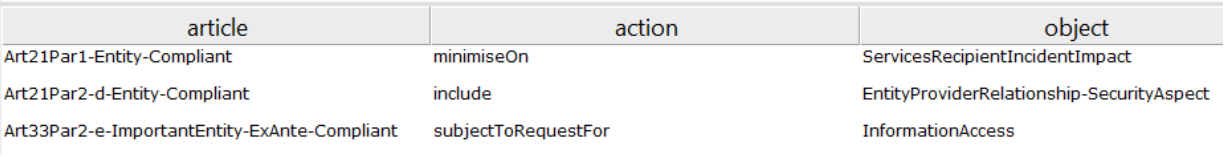

article	action	object
Art21Par1-Entity-Compliant	minimiseOn	ServicesRecipientIncidentImpact
Art21Par2-d-Entity-Compliant	include	EntityProviderRelationship-SecurityAspect
Art33Par2-e-ImportantEntity-ExAnte-Compliant	subjectToRequestFor	InformationAccess

Fig. 3. Result of the SPARQL Query

The SPARQL query computes a set subtraction between the measures specified by the compliance model, in our case, *ImportantEntity-ExAnte*, and the measures to which the provided entity is compliant. These measures represent the requirements that are missing from the company. Specifically, these are Article 12, paragraphs 1 and 2, and Article 33 paragraphs 2 (first column). The query also shows what are the actions to be taken (second column) and the objects to which the action should be applied (third column) to attain full compliance: this interaction increases the company's awareness on the tasks to carry out towards compliance.

5. Conclusions

In our paper, we presented NIS2Onto, a novel OWL ontology designed to model the NIS 2 Directive and the related security measures. Our work addresses the pressing need for a structured and interoperable framework that can effectively model legal and technical requirements, facilitating compliance and enhancing understanding across diverse stakeholders. NIS2Onto provides a comprehensive representation of the key concepts of the directive, including entities, relationships, and obligations, thereby offering a robust tool for compliance management. NIS2Onto enables the automation of compliance verification processes, supports risk assessment, and fosters improved communication among cybersecurity professionals, legal experts, and organisational leaders.

Our evaluation of NIS2Onto, using a combination of ontological metrics and qualitative analysis, demonstrates its effectiveness in encapsulating the requirements of the directive and its potential for practical application. While ontological metrics offer insights into the structural attributes of NIS2Onto, the contribution of the ontology is to facilitate nuanced interpretations of complex legal texts and support decision-making processes.

Future work will focus on extending NIS2Onto to incorporate additional cybersecurity frameworks and directives, enhancing its scalability and adaptability. We also aim to develop tools and interfaces that leverage NIS2Onto for automated compliance monitoring and reporting. Additionally, further research will explore the integration of NIS2Onto with other ontological systems, aiming to create a more unified and comprehensive approach to cybersecurity management. For this purpose, OASIS 2 can be used to describe in detail NIS 2 agents and their commitments [1, 3].

In conclusion, NIS2Onto represents a significant advancement in the field of ontologies for cybersecurity, offering an effective means of navigating the complexities of the NIS 2 Directive. It provides a foundation for ongoing research and development, paving the way for more sophisticated and integrated cybersecurity solutions.

[2] The SPARQL query is available at the repository provided.

Acknowledgments

Gianpietro Castiglione acknowledges a studentship by Intrapresa S.r.l. and Italian “Ministero dell’Università e della Ricerca” (D.M. n. 352/2022).

References

[1] G. Bella, G. Castiglione and D.F. Santamaria, A Behaviouristic Approach to Representing Processes and Procedures in the OASIS 2 Ontology, in: *Proceedings of the Joint Ontology Workshops 2023, Episode IX: The Quebec Summer of Ontology, co-located with the 13th International Conference on Formal Ontology in Information Systems (FOIS 2023), Sherbrooke, Québec, Canada, July 19–20, 2023*, Vol. 3637, CEUR Workshop Proceedings, 2023, pp. 1–17.

[2] G. Bella, G. Castiglione and D.F. Santamaria, An Ontological Approach to Compliance Verification of the NIS 2 Directive, in: *Proceedings of the Joint Ontology Workshops 2023, Episode IX: The Quebec Summer of Ontology, co-located with the 13th International Conference on Formal Ontology in Information Systems (FOIS 2023), Sherbrooke, Québec, Canada, July 19–20, 2023*, Vol. 3637, CEUR Workshop Proceedings, 2023, pp. 1–12.

[3] G. Bella, D. Cantone, C.F. Longo, M. Nicolosi-Asmundo and D.F. Santamaria, The ontology for agents, systems and integration of services: OASIS version 2, *Intelligenza Artificiale* **17**(1) (2023), 51–62. doi:10.3233/IA-230002.

[4] G. Castiglione, G. Bella and D.F. Santamaria, Towards Grammatical Tagging for the Legal Language of Cybersecurity, in: *Proceedings of the 18th International Conference on Availability, Reliability and Security*, ARES ’23, Association for Computing Machinery, New York, NY, USA, 2023. ISBN 9798400707728. doi:10.1145/3600160.3605069.

[5] G. Castiglione, G. Bella and D.F. Santamaria, Seconto: Ontological Representation of Security Directives, 2024, Available at SSRN: https://ssrn.com/abstract=4862271.

[6] L. Elluri, A. Nagar and K.P. Joshi, An Integrated Knowledge Graph to Automate GDPR and PCI DSS Compliance, in: *2018 IEEE International Conference on Big Data (Big Data)*, 2018, pp. 1266–1271.

[7] ENISA, NIS 2 Directive source, 2022. https://eur-lex.europa.eu/eli/dir/2022/2555/oj.

[8] S. Fenz and T. Neubauer, Ontology-based information security compliance determination and control selection on the example of ISO 27002, *Inf. Comput. Secur.* **26** (2018), 551–567. https://api.semanticscholar.org/CorpusID:54458814.

[9] M. Fernández-López, A. Gómez-Pérez and N. Juristo, METHONTOLOGY: From Ontological Art Towards Ontological Engineering, in: *Proceedings of the Ontological Engineering AAAI-97 Spring Symposium Series*, American Asociation for Artificial Intelligence, 1997, pp. 1–8, Ontology Engineering Group - OEG. https://oa.upm.es/5484/.

[10] G.B. Gianpietro Castiglione Daniele Francesco Santamaria, NIS2Onto. https://github.com/gianpietroc/nis-ontology.

[11] K.P. Joshi, A. Gupta, S. Mittal, C. Pearce, A. Joshi and T. Finin, Semantic approach to automating management of big data privacy policies, in: *2016 IEEE International Conference on Big Data (Big Data)*, 2016, pp. 482–491. doi:10.1109/BigData.2016.7840639.

[12] R. Matsunaga, I. Ricarte, T. Basso and R. Moraes, Towards an Ontology-Based Definition of Data Anonymization Policy for Cloud Computing and Big Data, in: *2017 47th Annual IEEE/IFIP International Conference on Dependable Systems and Networks Workshops (DSN-W)*, 2017, pp. 75–82. doi:10.1109/DSN-W.2017.28.

[13] J.H. Michael Poppe Martin Lichtwark, OntoMetrics — ontometrics.informatik.uni-rostock.de, howpublished = https://ontometrics.informatik.uni-rostock.de/ontologymetrics/.

[14] J. Muñoz, G. Esteban, O. Corcho and F. Serón, PPROC, an ontology for transparency in public procurement, *Semantic Web* **7** (2016), 295–309. doi:10.3233/SW-150195.

[15] M. Palmirani, M. Martoni, A. Rossi, C. Bartolini and L. Robaldo, PrOnto: Privacy Ontology for Legal Reasoning, in: *International Conference on Electronic Government and the Information Systems Perspective*, 2018. https://api.semanticscholar.org/CorpusID:52047214.

[16] J. Raad and C. Cruz, A Survey on Ontology Evaluation Methods, in: *International Conference on Knowledge Engineering and Ontology Development*, 2015. https://api.semanticscholar.org/CorpusID:27547076.

[17] R. Syed, Cybersecurity vulnerability management: A conceptual ontology and cyber intelligence alert system, *Inf. Manag.* **57** (2020), 103334. https://api.semanticscholar.org/CorpusID:225269944.

[18] Z. Syed, A. Padia, T.W. Finin, M.L. Mathews and A. Joshi, UCO: A Unified Cybersecurity Ontology, in: *AAAI Workshop: Artificial Intelligence for Cyber Security*, 2016. https://api.semanticscholar.org/CorpusID:6896947.

[19] L. Tailhardat, Y. Chabot and R. Troncy, NORIA-O: An Ontology for Anomaly Detection and Incident Management in ICT Systems, Springer-Verlag, Berlin, Heidelberg, 2024, pp. 21–39–. ISBN 978-3-031-60634-2. doi:10.1007/978-3-031-60635-9_2.

[20] J.A. Wang and M. Guo, OVM: an ontology for vulnerability management, in: *Cyber Security and Information Intelligence Research Workshop*, 2009. https://api.semanticscholar.org/CorpusID:3331211.