(Meta)data for Official Government Information: the TOOI Ontology and Knowledge Graph*

Jan Voskuil[1,*], Marc van Opijnen[2], Hans Overbeek[2], Theun Fleer[2] and Wessel Schollmeijer[2]

[1] *Taxonic B.V., The Netherlands*
[2] *KOOP, Ministry of Interior and Kingdom Relations, the Netherlands*

Abstract
The TOOI knowledge graph aims to achieve the FAIR objectives for official government information in the Netherlands. Its relevance extends beyond the immediate context in which it is conceived. This article presents the general characteristics of TOOI, how its constituting parts interrelate, and how its sustainability as a living standard is managed. It focuses on its core component, the TOOI ontology, and discusses some aspects of its design and development. It discusses how ODCM and OntoUML were applied, and reflects on practical aspects of the application of these methods.

Keywords
open government, (meta)data, ontoUML, ontology.

1. Introduction

TOOI [1] (acronym for 'Thesauri and Ontologies for Official Government Information') is a reference model in which authoritative information about public organizations and open government information is made available in a structured and machine-readable format for the purpose of coherence and findability of such information from various sources. Ultimately, TOOI's goal is to make such information FAIR [1]. This article focuses on the TOOI ontology in the context of the broader knowledge graph.

1.1. Problem statement

In today's complex and highly digitalized society, public transparency is of the utmost importance. Not only in complex crises, like the Covid-19 outbreak, but also in day-to-day life, lawyers, journalists, businesses, special interest groups and the general public at large, but also public organizations themselves, increasingly need coherent official documents and public data from a variety of sources. For instance, all those stakeholders have to be

FOIS 2024, Ontology showcase, July 15–19, 2024, Enschede, The Netherlands
* Part of the material in this paper is also available in Dutch as part of the TOOI documentation, see references.
* Corresponding author.
jan.voskuil@taxonic.com (J. Voskuil); {marc.opijnen, hans.overbeek, theun.fleer, wessel.schollmeijer} @koop.overheid.nl

informed about plans and decisions that impact their physical and economic environment. The following problems arise:

1. It is difficult to find and combine pieces of information when these are distributed over different official documents from a variety of sources;
2. Exchange and reuse of information between government organizations (interoperability) entails reconciliation of disparate metadata, which is costly;
3. Different organizations keep local registries of the same reference data using different conventions, causing incompatibilities (thus, the municipality Bergen in Limburg is variously referred to as Bergen (L), Bergen (L.), Bergen (Li), municipality of Bergen (L), etc.) and unnecessary costs.

The existence of these problems and the need to address them has been recognized in various policy documents of the Dutch government (e.g., [3]), as well as in academic papers [4]. The urgency increases since recent legislation prescribes that information from all types of public organizations must be actively published and made available in a coherent manner. The relevant legislation includes:

- the Open Government Act (Woo), replacing as of 1 May 2022 the Freedom of Information Act (Wob);
- the Promulgation Act (Bw), that has been amended twice in recent years as to make the electronic promulgation authentic and widen the scope of official documents to be disseminated electronically;
- the Environment and Planning Act (Ow);
- the Act on the re-use of public sector information (soon including the implementation of the EU Open Data Directive as well).

It has been noted that similar problems exist in other countries [5]. As we will see, TOOI, including its ontology, is based on the situation in the Netherlands. Details of legislation and government structure do not directly carry over to other countries. However, the methodology underpinning TOOI is applicable in any jurisdiction.

1.2. The added value of TOOI as a solution

A coherent view of information across public bodies requires a standardized language which clearly states (for humans and computers) which words or symbols are used for which things. This standard is TOOI, a knowledge graph representing certain aspects of the structure of the entire Dutch government, both central and regional, and of the structure of official, public information produced by public bodies. TOOI supplies standardized names for structural things (classes and properties) and for things used as reference data, such as identifiers for individual organizations. TOOI is created and maintained by KOOP, the publications office of the Netherlands (portal: https://www.koopoverheid.nl/). KOOP and others can use TOOI to express metadata.

To support large-scale publication (paragraph 1.1), KOOP provides a number of facilities. Government organizations must provide information via these central facilities, including prescribed metadata. These facilities are:

- The electronic promulgation of legislation and the publication of all official announcements from all public, at the national as well as the regional level. Portal: officielebekendmakingen.nl;
- The Digital infrastructure for the Open Government Act on which public information from most government organizations must be made available. Portal: open.overheid.nl;
- The standard for official publications and consolidated regulations (STOP) used in the Digital System for the Environment and Planning Act (DSO);
- The national facility for the publication of those STOP documents;
- The database with consolidated legislation of the central government. Portal: wetten.overheid.nl;
- The database with consolidated legislation of all regional authorities). Portal: lokaleregelgeving.overheid.nl;
- The Register of Government Organizations. Portal: organisaties.overheid.nl;
- The Data Register of the Dutch Government at data.overheid.nl.

As noted, inconsistencies in the metadata used by these facilities are problematic. In order to meet the requirements of the Woo, uniformity and quality of these metadata must be improved. To achieve this goal, the TOOI knowledge graph is increasingly applied.

1.3. From Organizations to Documents and Back

The starting point of TOOI is the realization that official publications and government organizations are interdependent. An official publication exists by virtue of being published by a government organization, as registered in its metadata. Conversely, a government organization comes into existence because of a legally grounded decision that comes into force only when it is published as an official publication. Thus, information objects are instrumental in causing government organizations to come into existence, to go out of existence, and to change their properties. The register of government organizations is also kept by KOOP (paragraph 1.2).

1.4. TOOI, the Knowledge Graph for Government Information

TOOI is a comprehensive knowledge graph expressed in RDF [7] consisting of:

- An ontology that specifies a domain language for describing organizations, information objects and their provenance;
- Registries that constitute the authentic government organization database;
- Thesauri that define concepts for classification purposes;
- Authority tables that are generated from the registries and thesauri.

Each of these modules is an RDF graph. Based on the RDF mechanisms for graph coalescence (e.g., owl:imports) these modules form a holistic, unified knowledge graph comprising all of them, while enabling users to easily select only the information they need. The authority tables further enhance this ease of selection. A TOOI authority table is essentially a set of triples extracted from either a thesaurus or a registry. It is then made available in various formats: Turtle, XML/RDF, JSON LD, and also plain XML.

The TOOI ontology (including the ontologies it includes, see paragraph 2) plays a central role in the knowledge graph, since it provides the structural conceptual model in terms of which the other modules (subgraphs) are expressed.

Figure 1 provides a visualization of the TOOI knowledge graph. Each filled rectangle signifies a separate module (RDF graph). The translucent grey rectangle on top contains the two graphs that together constitute the ontology. Beneath it on the left are the registries and thesauri. On the right we have the authority tables.

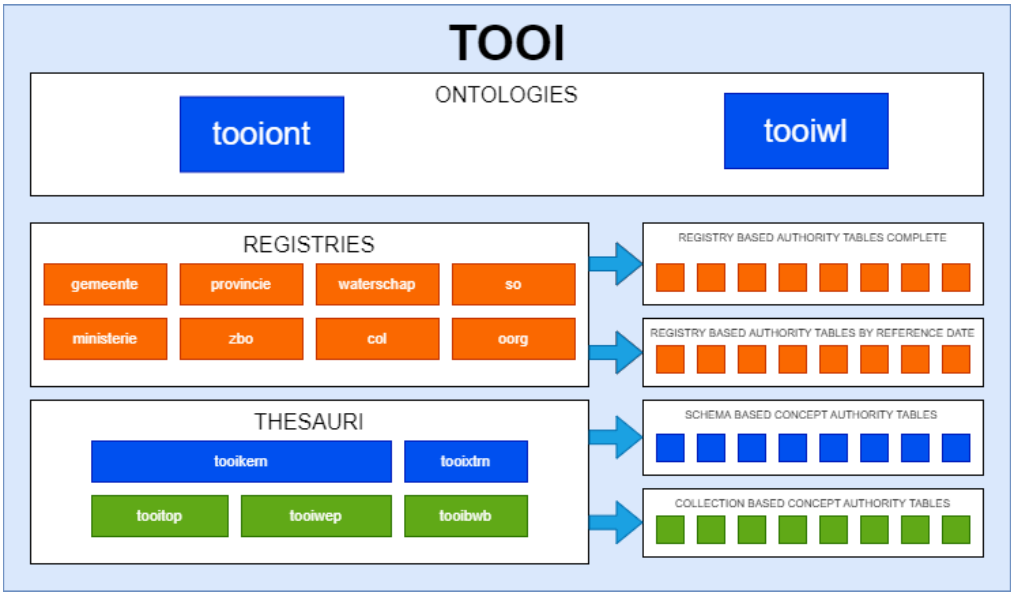

Figure 1: An overview of the TOOI modules constituting the TOOI knowledge graph.

2. The TOOI Ontology

Let us now zoom in on the ontology that underpins TOOI, the TOOI ontology [4]. The ontology is expressed in RDF. Some resources from RDFS [8] and OWL [9] are used. Constraints on properties (relations and attributes) are expressed in SHACL [10]. A UML-diagram is available as part of the documentation. The documentation of the ontology is a mix of hand-crafted text and diagrams, and data from the ontology's RDF-representation, which is automatically inserted using scripting.

The TOOI ontology makes pervasive use of other ontologies, including SKOS [11], Dublin Core [12], DCAT [13], ADMS [14], PROV-O [15], and ORG [17], among others.

2.1. The domain represented by the ontology

The ontology covers government organizations, information objects, and provenance.

2.1.1. Government organizations

Government organizations have many aspects. The TOOI ontology takes primarily a legal perspective: what does the law say about them? What it says is in large part highly precise and conducive to formalization. In other parts, aspects of organizations are grounded in history, culture and tradition. These aspects can be quite obscure. Still, some of these aspects need to be included as authentic information in the registries. We will return to this point in paragraph 2.2 below.

2.1.2. Information objects

TOOI defines information objects as objects of which all essential properties can be conveyed in a message, following [17]. In its present form, the TOOI ontology only touches on fairly generic points (which are still quite manifold). The class Information Object has only a limited number of subclasses. Properties defined for this class are of a relatively generic nature.

2.1.3. Provenance

The third area covered by the TOOI ontology is provenance, including organizational history. This concerns an extension of PROV-O. We distinguish between, on the one hand, existential change events, in which (for instance) an organization comes into existence or goes out of existence, and, on the other hand, change of state events, in which (for instance) the name of an organization changes. Such a change generates a representation of a static time-slice of the organization, which we call a historical version. Historical versions have their own identity. This treatment of history follows a pattern discussed in [18] (sections 4.4.1 and 6.5).

Historical versions are useful for, e.g., presenting metadata of documents to human users. For example, when a document was published by some organization at a certain time, it is relatively easy to retrieve and display the name of that organization at that time, even though the current name of the organization may have changed. Conversely, when looking for all publications of a certain organization, one can simply use the organization's identifier, ignoring any name changes. A more detailed discussion is provided in [19].

This ontological pattern for representing history has recently been applied in other ontologies as well.

2.2. Ontological design patterns

The TOOI ontology applies several common design patterns. One that is pervasively used is classification using SKOS concepts. TOOI uses the subclass relation in cases where the relation between superclass and subclass meets several or all of a list of criteria. These include: (i) the intention of the subclass can be clearly defined; (ii) the extension of the subclass at any point in time is tractable from a legal perspective; (iii) instances of the subclass have one or more properties that are unique for that subclass. Whenever these criteria cannot be sufficiently met, we use a custom classification property that takes instances of the class Concept (as defined in SKOS) as value.

For example, the law provides a clear definition of what a municipality is. The rules for creating new municipalities and "decommissioning" old ones are equally clear. When a new municipality is created, the decision is published in an official act that is, effectively, an integral part of Dutch legislation. Also, municipalities have the unique property that they are located in a province. Thus, the TOOI ontology defines a class Municipality which specializes the class Government Organization, and the extension of this class at any point in time can be retrieved from the TOOI registries.

The notion of a "adviescollege" (*advisory board*) is murkier. For these and other cases, definitions may be less clear or absent, and in determining their extension, a certain level of subjectivity may be involved. To deal with such cases, TOOI makes use of SKOS and defines a concept with this label. The concept is then used to categorize the relevant organizations. Technically, these objects have a property called *organizationKind* which, in the appropriate cases, takes said concept as value. This pattern for expressing classification contrasts with the subclassing pattern by avoiding any specific ontological commitment. The TOOI ontology makes extensive use of it. In the Appendix, we analyze this pattern in terms of OntoUML.

2.3. Other ontologies

There are various international standards for describing and/or marking-up legislation and its constituting parts, such as Akoma Ntoso [20] and ELI [21]. The TOOI ontology presently stays clear of modeling in any detail the specifics of acts, policy guidelines, regulations and other forms of legislation.

Another ontology (or, more specifically, metadata schema) that is relevant in the domain of metadata for official information in the Netherlands is MDTO [22]. MDTO (Metadata for sustainably accessible government information) is a schema for recording and exchanging metadata to enable sustainable accessibility of government information. The teams behind MDTO and TOOI collaborate on detailing the relation between the two. This is ongoing work.

2.4. Methodology: UFO and OntoUML

Implicitly, use is made of UFO and OntoUML [18][23][24][25]. This enables us to keep track of overall quality. At the same time, because UFO captures ontological patterns in terms of meta-universals, it is easy to hide the technical apparatus from end users. To support uptake, it is important to stay close to the experience of users. Therefore, in UML-diagrams that are part of the TOOI documentation we do not use OntoUML stereotypes.

OntoUML is used internally to guarantee formal correctness. Thus, whenever classes and properties are added, OntoUML stereotypes are applied. When this leads to the discovery of issues, these are resolved. This is not to say that application of OntoUML is always uncomplicated: sometimes dilemmas arise. In the Appendix, a partly decorated class diagram of the current version of the TOOI ontology (1.4.0) is presented, along with a discussion of some of these questions. In paragraph 3, we reflect on the importance of OntoUML for creating models like TOOI and how further uptake could be stimulated.

2.5. Requirements driving the development of the TOOI ontology

The requirements driving the development of the TOOI ontology are subordinate to the problems that the TOOI knowledge graph as a whole intends to solve (paragraph 1.1). The more specific requirements for the ontology are:

- Define the structural aspects of the domain that TOOI as a whole describes;
- Provide consistent mechanisms for dealing with temporality and change;
- Support common understanding and communication, and afford psychological groundedness [26];
- Support relatively simple approaches (queries) for common retrieval tasks;
- Provide ample documentation to facilitate uptake;
- Reuse ontologies that are widely used in the domain where possible.

2.6. How and where the TOOI ontology is used

Presently, the TOOI ontology is mainly used inside the TOOI knowledge graph itself. The identifiers of organizations and other entities are increasingly used in metadata. This gradually promotes use of the ontology resources as well. Reusing these resources can be done indirectly, by using custom XML- or JSON-elements that are (implicitly or explicitly) mapped to these resources, as is the case in the portal for the Open government act (paragraph 1.2). Direct use of TOOI ontology resources through, for instance, JSON LD APIs, will be possible in the near future, when the first APIs will be made available to directly interact with the knowledge graph.

Thus, in the current situation, the TOOI ontology is primarily used inside of the KOOP organization, specifically, in the publication platforms listed in paragraph 1.2 — mostly indirectly. In the longer term, it is expected that the ontology will be increasingly used directly, inside and outside of KOOP.

2.7. Evaluating the ontology and acting on the results

The structural quality aspects of the TOOI ontology are evaluated by applying the OntoUML methodology, as described in paragraph 2.4. The usability aspects are evaluated by applying the ontology in practical settings.

2.7.1. Evaluating usability

The TOOI ontology is developed incrementally. At each increment, datasets are created that use the ontology for expressing information, such as organization registries. By using the ontology in such datasets, its usability aspects can be evaluated, including understandability and effectiveness, as well as the way these datasets can be queried.

The use that third parties make of the TOOI authority tables is difficult to monitor, but through interaction with known users, further insight is gained. Such insights may lead to changes in the ontology or in guidance on how it should be used.

2.7.2. Acting on the results: an example

As an example, we discuss a policy change that is currently underway. Recall that when a change occurs in the state of an organization, a historical version of that organization is registered (paragraph 2.1.3). The historical version has its own identity and persistent identifier. It represents the static, unchangeable state that the organization was temporarily in. The state is also registered at the "organization level." The state of the organization, that is, the totality of its property values, changes through time. Unlike the states of its historical versions, the state of the organization is dynamic. Ontologically, historical versions are weak entities in the sense that their identity is not preserved across possible worlds. Direct reference to these entities is best avoided also for practical reasons.

If a developer needs to present metadata of an information object in a human-readable form, they typically need to access the name that the organization had at the moment it published that information object — not its present name, if it has changed. To achieve this, the developer should query the knowledge graph, retrieve the correct historical version, and take the organization's name from there.

However, based on feedback received, we noticed that many developers would prefer to include a reference to the historical version of the organization in the metadata of the information object. There is an obvious danger in allowing this. For example, if the developer would register the historical version as the publisher, two problems arise. First, it is ontologically wrong. Organizations publish, not their historical versions. Second, it would run afoul of the FAIR principles. Two information objects published by the same organization at different moments will not be recognized as such when different identifiers are used to refer to (the historical versions of) the organization.

A pragmatic solution would be for the developer to register in the metadata, alongside the identifier of the organization, the pertinent name as a string. Developers, however, do not like this solution. The problem is also solved if in the metadata the developer registers the organization as the publisher, *and* the relevant historical version as the value of some custom metadata field. Until now, measures have been in place that make it impractical for developers to directly refer to historical versions. In the near future, we intend to backtrack this policy decision. To avoid errors, specific implementation guidelines will be provided.

The constant evaluation of the ontology (including its use) and the implementation of interventions as a result is a continuous process. This process is currently being formalized as described below in paragraph 2.9.

2.8. TOOI and FAIR

The TOOI knowledge graph is modular and consists of separate RDF. Each of these, including the ontology, complies with the FAIR principles. The management documentation describes in detail how each module implements each of the FAIR principles [27].

2.9. Sustainability

TOOI's sustainability policy implements and complies to BOMOS [28]. BOMOS is a Dutch standard for managing, maintaining and developing open standards. A full description of TOOI's sustainability policy is provided in its management documentation [27]. Here, we present some highlights. Note that parts of the following description are currently still in the process of being implemented.

2.9.1. Version policy

TOOI is a living standard and will keep evolving. Change is necessary to keep the content of TOOI in sync with the changing reality, to implement evolving insights and to provide new features needed by users. The constant need for change can be at odds with the need for persistence required for interoperability of applications that use TOOI. The solution is a carefully defined change process and transparent version management of all TOOI modules.

TOOI is developed in an agile and iterative manner. New versions of each module will be released regularly. TOOI follows the principles of semantic versioning. Versions, once published, remain available as long as TOOI is supported. Historical and development versions will be available through GitLab.

2.9.2. Change management

Changes can be initiated by anyone by submitting issues in one of the TOOI modules in GitLab, or via e-mail. The editors process the request and keep the submitter informed of progress. Plans for modifying TOOI specifications will be coordinated in the open TOOI user consultation group (currently being instituted). If accepted, the modifications become part of the published development roadmap.

2.9.3. The change process

In the change process, a change request goes through the phases of content evaluation, assessment and aesting, decision making and implementation. Substantive correctness is determined during testing, and the desirability of the proposed changes is evaluated during decision-making. The result of the Assessment and Testing phase is a (possibly adapted) change request for the module.

3. Reflection: Ontology Driven Conceptual Modeling in practice

The TOOI knowledge graph is a complex whole. The design process must balance many conflicting requirements. This is certainly true for TOOI's core artifact: the ontology. Using methodologies from the field of Ontology Driven Conceptual Modeling (ODCM) in general, and, UFO and OntoUML in particular, has helped the design process. Specifically, it has provided guidance and depth to discussions on class hierarchy and relations, and on defining the right concepts in the right way.

This is in line with results reported in [29], which show that OntoUML significantly contributes to improving the quality of structural conceptual models without requiring an additional effort to producing them. It is not surprising, therefore, that, in comparison to other ODCM approaches, OntoUML is relatively widely used, as discussed in [25].

Conversely, however, a majority of domain modeling projects in the public sphere do not use ODCM methods at all, as even a cursory inventory will show. ODCM is underused, which leads to models that, despite the best of intentions, fail to produce the intended level of semantic interoperability. Without a solid ontological practice, achieving the FAIR objectives will remain elusive [30].

The only way to improve this situation is to encourage a higher adoption rate of ODCM. What can we, as a community, contribute? We would like to mention three specific factors that play a role. First, ODCM is a vibrant field of research, and fascinating results are reported in a growing stream of research. These could be translated more proactively into practical guidelines. A naïve web search turns up tools and websites on OntoUML based on OntoUML 1.0, whereas OntoUML 2.0 dates back to 2018 [24]. Are the insights offered in [31] on the ontology of events part of OntoUML 2.0 or not?

This latter question brings us to the second factor. What is needed is a more formal description of OntoUML — with 'formal' in the sense of a formal standard. This is not so much about getting it declared as a standard at, say, W3C, OGC or ISO. Rather, an adequately funded organization is called for that structurally provides a workable level of sustainability, including clarity of scope and versioning.

The third factor is education. Most of the materials on UFO and OntoUML are available as academic journal articles, doctoral dissertations and the like. For practicing information analysts in the industry, this is not easily accessible. Commercial course materials on UFO specifically oriented towards such an audience start becoming available [32].

Thus, we advocate more practical descriptions, a sustainable standardization effort, and more educational resources. Progress on these fronts will, hopefully, foster further uptake of OntoUML, so that more projects can leverage the benefits of ODCM, bringing us all closer to achieving the goals of FAIR Data.

References

[1] KOOP, TOOI, 2024. KOOP standard. URL: https://standaarden.overheid.nl/tooi

[2] M. Wilkinson, M. Dumontier, I. Aalbersberg, The FAIR Guiding Principles for scientific data management and stewardship, Sci Data 3 (2016).

[3] Ministry of the Interior and Kingdom Relations, Geactualiseerde Werkagenda Waardegedreven Digitaliseren, 2024. URL: https://open.overheid.nl/documenten/8fb16ed3-0946-49d5-bf1a-96724f1762d6/file

[4] M. Marx, M. Larooij, F. Perasedillo, J. Kamps, Enticing Local Governments to Produce FAIR Freedom of Information Act Dossiers. In: Kamps, J., *et al.* Advances in Information Retrieval, ECIR 2023, Lecture Notes in Computer Science, vol 13982 (2023). doi: 10.1007/978-3-031-28241-6_25

[5] Lisa DeLuca, Searching FOIA Libraries for government information, Government Information Quarterly, Volume 37 (2020). doi: 10.1016/j.giq.2019.101417

[6] KOOP, TOOI-Ontologie, 2024. KOOP standard.
URL (documentation): https://standaarden.overheid.nl/tooi/doc/tooi-ontologie/
URL (source, Turtle): https://identifier.overheid.nl/tooi/def/ont.ttl

[7] Richard Cyganiak, David Wood, Markus Lanthaler, RDF 1.1 Concepts and Abstract Syntax, 2014. W3C Recommendation. URL: https://www.w3.org/TR/rdf11-concepts/

[8] Dan Brickley, Ramanathan Guha, RDF Schema 1.1, 2014. W3C Recommendation. URL: https://www.w3.org/TR/rdf-schema
[9] Boris Motik, Peter F. Patel-Schneider, Bijan Parsia, OWL 2 Web Ontology Language Structural Specification and Functional-Style Syntax (Second Edition), 2012. W3C Recommendation. URL: https://www.w3.org/TR/owl2-syntax
[10] Holger Knublauch, Dimitris Kontokostas, Shapes Constraint Language (SHACL), 2017. W3C Recommendation. URL: https://www.w3.org/TR/shacl
[11] Alistair Miles, Sean Bechhofer, SKOS Simple Knowledge Organization System Reference, 2009. W3C Recommendation. URL: https://www.w3.org/TR/skos-reference
[12] DCMI, Dublin Core Metadata Element Set, Version 1.1, 2012. DCMI Recommendation. URL: http://dublincore.org/documents/dces/
[13] Riccardo Albertoni, David Browning, Simon Cox, Alejandra Gonzalez Beltran, Andrea Perego, Peter Winstanley, Data Catalog Vocabulary (DCAT) - Version 2, 2020. W3C Recommendation. URL: https://www.w3.org/TR/vocab-dcat-2/
[14] Phil Archer, Gofran Shukair, Asset Description Metadata Schema (ADMS), 2013. W3C Recommendation. URL: https://www.w3.org/TR/vocab-adms/
[15] Timothy Lebo, Satya Sahoo, Deborah McGuinness, PROV-O: The PROV Ontology, 2013. W3C Recommendation. URL: https://www.w3.org/TR/prov-o
[16] Dave Reynolds, The Organization Ontology, 2014. W3C Recommendation. URL: https://www.w3.org/TR/vocab-org/
[17] Ian Jacobs, Norman Walsh, Architecture of the World Wide Web, Volume One, 2004. W3C Recommendation. URL: https://www.w3.org/TR/webarch
[18] Giancarlo Guizzardi, Ontological Foundations for Structural Conceptual Models, Ph.D. thesis, CTIT, Centre for Telematics and Information Technology, 2005. URL: https://ris.utwente.nl/ws/portalfiles/portal/6042428/thesis_Guizzardi.pdf
[19] Jan Voskuil, Marc van Opijnen, Hans Overbeek, History in the TOOI knowledge graph, 2022. URL: https://www.linkedin.com/pulse/history-tooi-knowledge-graph-jan-voskuil
[20] Akoma Ntoso, 2024. OASIS standard. URL: https://standict.eu/standards-repository/standard/akoma-ntoso-version-10
[21] Publications office of the European Union, ELI, 2024. Eur-lex standard. URL: https://eur-lex.europa.eu/eli-register/resources.html
[22] Nationaal Archief, MDTO. URL: https://www.nationaalarchief.nl/archiveren/mdto
[23] Giancarlo Guizzardi, G. Wagner, J.P.A.A. Almeida, R.S. Guizzardi, Towards ontological foundations for conceptual modeling: the unified foundational ontology (ufo) story, Applied ontology 10(3-4), 259–271 (2015)
[24] Giancarlo Guizzardi, Claudenir M. Fonseca, Alessander Botti Benevides, João Paulo A. Almeida, Daniele Porello, Tiago Prince Sales, Endurant Types in Ontology-Driven Conceptual Modeling: Towards OntoUML 2.0, in: Proceedings of the 37th International Conference on Conceptual Modeling (ER 2018), 2018. doi: 10.1007/978-3-030-00847-5_12

[25] Giancarlo Guizzardi, Alessander Benevides, Claudenir Fonseca, Daniele Porello, João Almeida, Tiago Prince Sales, UFO: Unified Foundational Ontology, Applied Ontology 1 (2021). doi: 10.3233/AO-210256
[26] John Mylopoulos, Conceptual Modelling and Telos, Tech. Report DKBS-TR-91-3, U. of Toronto (1991). URL: https://api.semanticscholar.org/CorpusID:422646
[27] KOOP, TOOI – Beheerplan, 2024. KOOP standard. URL: https://standaarden.overheid.nl/tooi/doc/tooi-beheerplan/
[28] Erwin Folmer, Gül Işik, Edwin Wisse, BOMOS, het fundament 3.0.1, 2023. Logius standard. URL: https://gitdocumentatie.logius.nl/publicatie/bomos/fundament/
[29] M. Verdonck, F. Gailly, R. Pergl, G. Guizzardi, B. Martins, O. Pastor, Comparing traditional conceptual modeling with ontology-driven conceptual modeling: an empirical study, Information Systems, 81:92–103 (2019).
[30] Giancarlo Guizzardi, Ontology, ontologies and the "I" of FAIR, Data Intelligence, 2 (1-2): 181–191 (2020).
[31] Nicola Guarino, Riccardo Baratella, Giancarlo Guizzardi, Events, their names, and their synchronic structure, Applied ontology 17 (2): 249-283 (2022).
[32] Jan Voskuil, Launching a course on Conceptual Modeling and UFO, 2024. URL: https://www.linkedin.com/pulse/conceptual-modeling-ufo-basics-taxonic-oajye/
[33] Nicola Guarino, Giancarlo Guizzardi, Processes as variable embodiments, Synthese 203 (4): 1-27 (2024).
[34] A. Albuquerque, Giancarlo Guizzardi, An ontological foundation for conceptual modeling datatypes based on semantic reference spaces, IEEE 7th International Conference on Research Challenges in Information Science (RCIS), 1–12 (2013).

A. Appendix: the TOOI ontology and OntoUML

The diagram in Figure 3 is taken from the TOOI ontology documentation [6]. Some labels have been provisionally translated into English (but note that legal concepts are not normally considered fully translatable). We added OntoUML 2.0 stereotypes as defined in [24]. These stereotypes are not part of the TOOI standard and are used only internally. Also, they are not intended to defend particular modeling choices, but rather to show how OntoUML helps in structuring design options and to elucidate some questions that arise.

As can be seen, in terms of UFO, the ontology's expression in OntoUML leaves many ontology elements implicit. In terms of [18], it suffers from *construct deficit*. This is a deliberate design choice. For example, most of the material relations lack explication of their truth makers. A relator kind is only introduced as a class in case we foresee a need to register information about its instances. Even in cases where relator kinds are explicitly introduced (for instance, the class MembershipCA, in blue near the bottom of the diagram), only parts of the full OntoUML notational apparatus are used. The list continues.

In all of such cases, the desideratum of notational completeness was weighed against useability requirements. Full notational completeness leads to a significant increase in the number of explicitly defined classes. Keeping the number small, on the other hand, improves (intuitive) understandability and ease in application. Moreover, whenever reference to specific implicit model elements is necessary for conceptual clarification or meaning negotiation, the missing notation elements can be transparently overlaid. Should practical problems arise, then explicit addition to the model notation can be achieved through the standard change process described in paragraph 2.9.3 above.

Let us now turn to questions that arise when applying OntoUML. Some such questions pertain to the exact boundary one draws in larger class hierarchies between categories and kinds. For example, in [33], it is observed that *killing* in *the killing of Ceasar* is an event category, whereas *stabbing* in *the stabbing of Ceasar* is an event kind: the manner of killing, namely, stabbing, introduces an identity principle. The practical question now is: is a change event as conceptualized in the diagram an event category? The diagram says it is. Further reflection on this question may reveal options for making the ontology more precise.

Another set of questions concerns the treatment of enumerations. In the context of RDF, enumerations are often encoded in terms of SKOS (paragraph 2.2). Can one treat this pattern in a unified way using OntoUML? In Figure 2 we have made an attempt to do so. In [34], Nominal Qualities are defined as quality universals the qualia of which originate from social conventions. This seems to capture the essence of what modelers try to express using SKOS concepts in structural conceptual models. We tentatively assign Concept to the metaclass ⟪kind⟫. In the *underlying* model, a subclass of Concept is the specific enumeration as provided in an authority table (paragraph 1.4) which, in turn, is based on a specific concept scheme or concept collection. The diagram displays a small selection of (currently 92) options. The same treatment applies to every other concept-taking property in Figure 3.

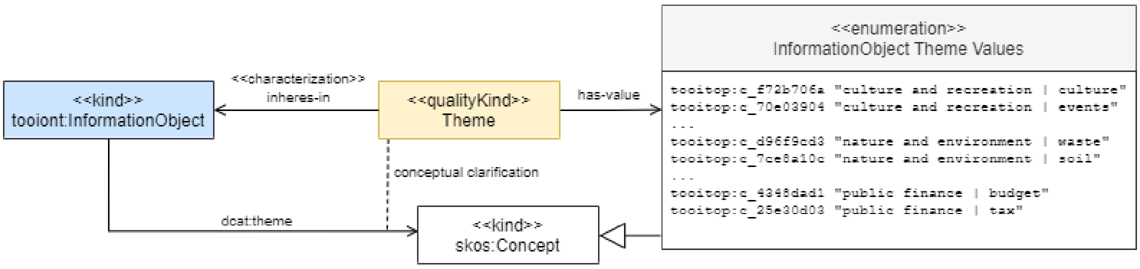

Figure 2: Conceptual clarification of how concepts (SKOS) are used in the ontology

Making explicit (in a machine-readable way) for each of these properties exactly which enumeration is applicable in which context is one of the larger tasks on the TOOI backlog. Part of how we intend to achieve this is by defining a subclass of Concept for each case, and leverage these subclasses in property shapes (SHACL) to describe the pertinent constraints. Practical benefits aside, this will cause the expression of the ontology to reach a higher level of isomorphism with the underlying model.

A deeper analysis of Concept would take into account the fact that a concept can be part of multiple enumerations, and be added to and removed from enumerations through time, suggesting that enumeration classes behave like roles with respect to qualities. This is left for another occasion.

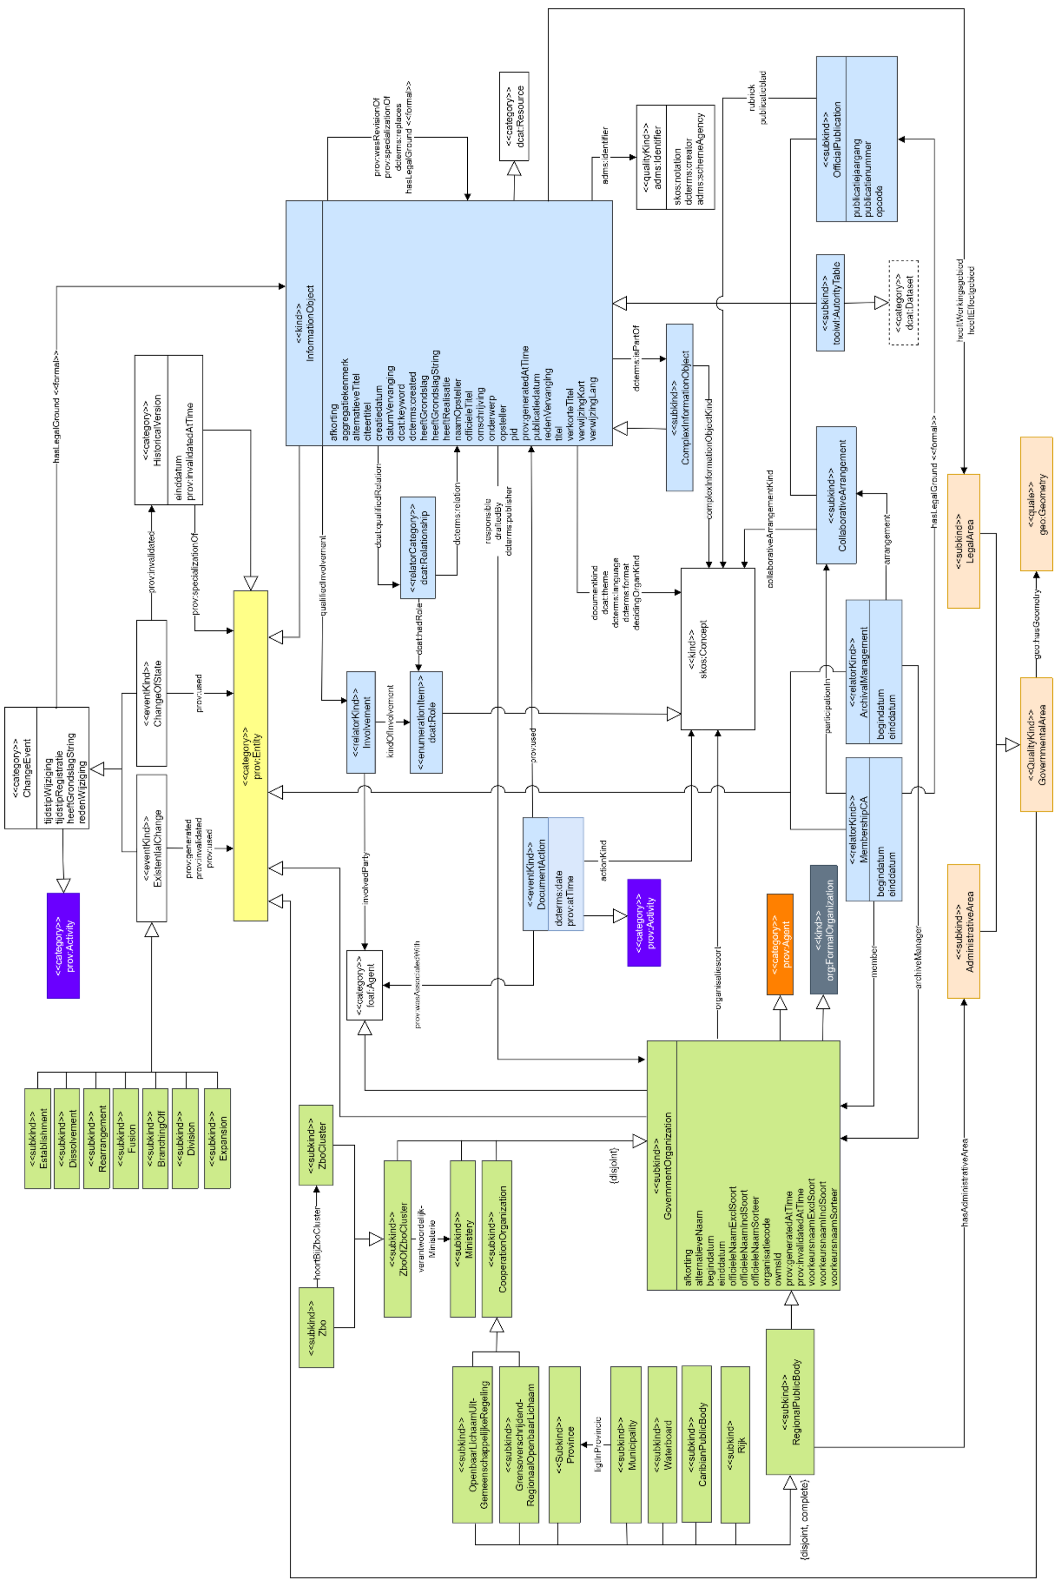

Figure 3: The TOOI ontology expressed in OntoUML (see [6])