Toward Requirements for an Ontology of Asset Management

Megan Katsumi[1,*], Tony Huang[2] and Mark S. Fox[1]

[1]*Department of Mechanical & Industrial Engineering University of Toronto, 5 King's College Rd. Toronto, ON M5S 3G8, Canada*
[2]*Toronto Water, 55 John St. Toronto, ON M5V 0C4, Canada*

Abstract
Asset management is comprised of data-intensive activities such as condition and performance monitoring, scheduling of repairs and preventive maintenance, capital project monitoring, and life-cycle costing. Like so many areas of industry, asset management is hindered by an evolution of legacy systems that has resulted in a complex web of disparate data systems making it difficult to track and report on assets. We aim to address this challenge for Toronto Water with the implementation of an ontology-supported data hub. Toronto Water is a major organization responsible for the treatment and supply of safe drinking water, the collection and treatment of wastewater, and stormwater management for the City of Toronto – serving a population of over 6 million people. In this paper, we focus on the identification of requirements for an ontology of asset management. This will be a key deliverable of the project that will serve to guide the survey of existing ontologies, as well as the development of the data hub itself. It is our hope that these requirements will serve to inform both the ontology and industry communities. The requirements not only identify the required scope for the ontology, but also illustrate the way in which ontologies can be used to support various asset management activities.

Keywords
asset management, ontology, requirements specification, water and sanitation

1. Introduction

Toronto Water is an organization responsible for providing services include delivering clean, safe drinking water, treating wastewater, and managing stormwater. Toronto Water provides 435 billion litres of drinking water and treats approximately the same amount of wastewater annually for more than 6 million residents and businesses within and around Toronto, Canada. In 2022, Toronto Water's budget consists of an operating budget of $1.4 billion CAD, including a $471 million component in capital reserve contribution.

Toronto Water's infrastructure has a replacement value of $83 billion. Due to pressures of population growth, infrastructure aging, and an increase in the scale of climate events, Toronto Water is currently in a decade-long period of infrastructure renewal. As the major asset changes come about, a sizeable challenge for Toronto Water's Asset Strategy and Maintenance Planning (ASMP) group is to track these changes in the various data systems. Recent legislation has also mandated that Toronto Water provide asset management plans to the provincial government.

FOMI 2022: 12th International Workshop on Formal Ontologies meet Industry, September 12-15, 2022, Tarbes, France
*Corresponding author.
katsumi@mie.utoronto.ca (M. Katsumi)

Asset management is a relatively new discipline. The common global understanding of asset management was gradually established with the publication of a number of authoritative standards and framework documents by the British Standard Institution [1], the Institute of Asset Management [2], the Asset Management Council, [3], and the International Standard Organization [4, 5]. These documents have all been published within the last 20 years. The staff of an organization starting to develop an asset management program might understandably equate it to its existing maintenance and reliability program, however the scope of an asset management however is much greater. Asset management enables an organization to examine the need for, and performance of, assets and asset systems at different levels. Additionally, it enables the application of analytical approaches towards managing an asset over the different stages of its life cycle.

The tasks involved in asset management activities are data-driven and thus are greatly hindered by the presence of data silos, as is the case at Toronto Water. In the ASMP group, efforts are under way to address this challenge with the development and implementation of a data hub. Central to the data hub design is an asset management ontology that will be used to facilitate data entry and access, and to enable more effective use of available data. Before proceeding further in the development process, it will be important to specify our requirements in greater detail. With an explicit understanding of the requirements, we will be able to more effectively assess existing work and better understand whether it can be reused, how it might fit together, and what development work may remain, in order to satisfy our needs.

This paper presents our progress toward the identification of our requirements for an asset management ontology; the remainder is structured as follows: we first provide background on the subject of asset management, and briefly discuss related ontologies. We introduce the approach taken to specify requirements for the ontology, then consider three key use cases and their associated requirements. We conclude with a summary of the required domains and thoughts on our next steps.

2. Background

In this section, we provide some background on the subject of asset management, discuss some existing ontologies for asset management, and introduce the methodology that will be adopted for the specification of requirements.

2.1. Asset Management

Asset management, broadly characterized, is a set of activities that gives an organization the knowledge and tools to fully utilize its physical assets to achieve its purpose [2]. An example of such an asset management activity is the formulation of a decision on whether to build a new piece of infrastructure, refurbish the existing one, or enhance its maintenance. Another example of an asset management activity is to make a future projection on the scale and placement of a particular technological capability needed for the near future - for instance, water pumping to serve a growing section of a city. Such projection is instrumental in planning for the capital funding of a project.

We can encapsulate Toronto Water's asset management objectives into the following statement: to maintain a high level of water quality and availability at a minimal cost while not compromising on the staff's exposure to any risk of injury or the public and the environment's exposure to harm.

It is helpful to consider achieving the objectives as a loosely defined optimization problem that requires a couple of categories of information to be approached methodologically. The first is the clarification of boundary constraints, i.e.

- the projected level of customer demand (and an inferred level of asset system capability to meet the demand);
- the expected product quality;
- the expected level of functional reliability and availability; and
- dangerous states of operation, deemed by environmental and safety regulations to be sufficiently likely to be harmful.

The second is the meaningful representation of relevant infrastructure; for example,

- the present condition of assets (many support the customer's primary needs, while some prevent the manifestation of injury or harm);
- the current age of assets, the condition and demand under which they operate, and their expected remaining useful life;
- the importance of an asset to immediate or long-term availability of an infrastructural capability that supports some quality of the product; and
- the likely impact of certain maintenance, refurbishment, or other actions on sustaining the asset's capabilities, as well as the budget.

The scope and complexity of information required to make defensible asset management decisions are far-reaching, and many pieces of input information must be derived from other measurable qualities or countable events. For example, information must be collected on the expected level of demand (for Toronto Water, this is tied to population forecast), environmental stresses, asset failure history, current asset condition, costs incurred to maintain the asset, and so on. Pieces of information such as assessments of criticality and condition are complex in that they are not directly measurable but rather the output of models. For example, the condition of an asset is inferred from certain measurable symptoms, such as the vibration amplitude of a specific machine part and the frequency of repair work seen in work history.

Information management is one of the primary challenges of asset management. Before even approaching the optimization problem described above, hundreds of variables and statements must be assigned to their context within a hierarchical semantic structure that reaches upward to the (seemingly abstract) objectives and downward to operational data streams.

2.2. Related Ontologies

Efforts have been made to develop asset management ontologies, however it's not yet clear whether or to what extent existing work may be reused for our project. Examples of such related work include the Tangible Capital Asset Ontology [6] which focuses solely on describing assets according to four different “modalities" such as the sector in which it is used, its aggregation level in the overall system (e.g., sub-system or component), and the type of function it performs. There are also maintenance ontologies, such as ROMAIN [7] and the ongoing work by the Industrial Ontologies Foundry (IOF) [8], specifically the Maintenance Working Group [9] which focuses on maintenance activities and failure states, but also naturally includes the concepts of asset, asset function, and role.

There is no de facto standard ontology for asset management and related ontologies vary in scope and purpose, so the decision of which ontology(s) to reuse, if any, is not straightforward. The purpose of this paper is not to provide an exhaustive survey of these ontologies, but to begin the identification of requirements which in turn will serve to inform the assessment and selection of existing ontologies.

2.3. ORSD Methodology

The Ontology Requirements Specification Document (ORSD) Methodology, presented in [10] was developed to provide guidance for the task of requirements specification. The method prescribes a series of steps to identify the requirements for an ontology, along with a template that can be used to specify the resulting requirements. After identifying the purpose, scope, and implementation language (Task 1); the end users (Task 2); and the uses (Task 3); the requirements are identified, grouped, validated, and prioritized (Tasks 4-7). Then, an initial set of required terms is identified based on the set of requirements (Task 8).

The specification of requirements is based on the use of Competency Questions (CQs) [11], complimented by the identification of what the authors label as ‘non-functional requirements’, drawing an analogy to software engineering. Rather than approaching the task of requirements specification in an ad-hoc way, the ORSD methodology imparts structure upon the approach and results.

3. Requirements

Currently, asset management employees are tasked with deciphering numerous data systems and schema in order to access the information they require. This task is exacerbated by the fact that many of the schema are not intuitive and are afflicted with idiosyncrasies owing to years of evolution and “layering-on" of features. As a result, asset management employees must often rely on I.T. professionals to parse and reconcile information from various systems in order to make use of it.

The goal of this project is to support asset management activities with an ontology-enabled data hub. The data hub will integrate the key data systems used by Toronto Water while also providing a single point of access to the data, as illustrated in Figure 1. The first benefit that will be realized from this integration is the removal of barriers for staff to access data. Two factors contribute to this. First, all relevant (available) information will be integrated via the data hub’s connection to the key data systems. This creates a single point of access to the data and eliminates the need for employees to search through multiple systems to retrieve the information they need. Second, the ontology will not only integrate the required data but

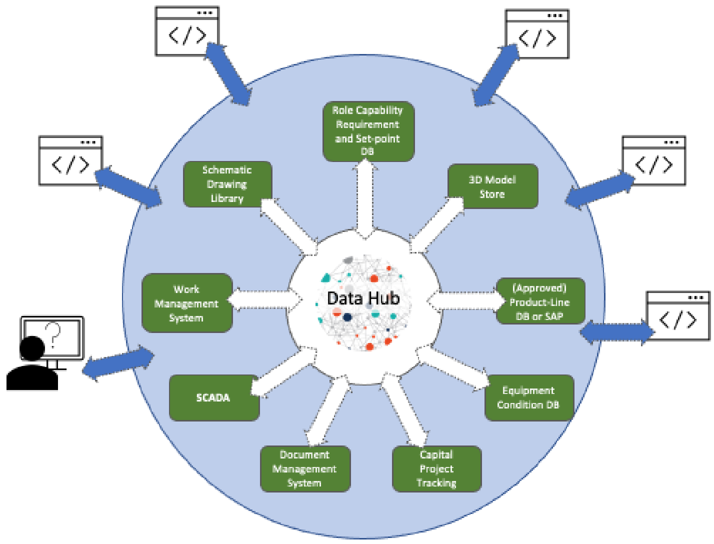

Figure 1: The Data Hub will integrate all key systems and serve as a single point of access for question answering while providing the data needed for special-purpose asset management applications.

provide an intuitive, standardized language for employees to access information (i.e., to pose and interpret queries).

Beyond this, the data hub is being developed to support more advanced functionality that will operate on the integrated data. This has been broken down into several development initiatives, each with its own set of use cases defined by the ASMP group at Toronto Water: historical data hub, content metadata, and data entry. In turn, each initiative is intended to enable a set of user-facing applications that will support specific asset management activities within Toronto Water. Each initiative is explained in greater detail below. Currently, development efforts are focused solely on the Historical Data Hub Initiative. Subsequent iterations will look at incorporating the other initiatives and applications. The ontology design effort will first tackle the creation of a foundational ontology for asset management, and then look at extending this representation with application-specific concepts (e.g. to capture the specific assets and processes that exist at Toronto Water).

The ORSD methodology described in the previous section has been adopted to specify the requirements for an asset management ontology for our project. These requirements will serve as a first step toward the development of an ontology for asset management. In this section, we present a detailed ORSD for the historical data hub initiative and preliminary ORSDs for the other initiatives that will be considered more comprehensively at later stages of development.

The following adaptations were made to the ORSD methodology:

- We omitted the grouping and prioritization of requirements. In fact, prioritization is identified as an optional task in the methodology. We may revisit this at later stages in development, however for now we assume that all of the requirements have equal priority, and there are not so many as to warrant a division into groups.

- We have not included terms from answers and objects in the pre-glossary at this point as these will point to Toronto Water-specific instances and classes that are not the focus of this stage of development. This will be the focus of later work when we extend this core representation to capture the specific assets and roles that exist at Toronto Water.

In each of the ORSDs we currently specify OWL as the required implementation language. This is largely motivated by a view of OWL as the de facto language for practical ontology implementation, and the view that it offers the greatest pool of, potentially reusable, existing ontologies. That is to say, it is not a firm constraint but a starting point for the specification.

3.1. Historical Data Hub Initiative

Asset history is built on accurate, precise, and complete records of asset change. No matter in which data system a particular change to an asset is captured, there is a need to describe the asset after the change as well as provide information about the change itself - for example, when, by whom, and in which system was the change captured? The ability to look back at an asset's history is an instrumental part of any asset management program. It is only through reliable tracking of the individual asset changes that a reliable life cycle history of the assets may be established. The life cycle history is foundational for estimating risk of failure, which informs (or should inform) most asset management decisions, including the advice that is provided to operations.

At Toronto Water, asset changes may be documented in different data systems including drawings, the work management system, geospatial information systems, spreadsheets, and internally developed applications. It is important to keep track of any discrepancies between systems and reconcile them routinely and timely, before divergent representations emerge in the systems.

The Historical Data Dub Initiative focuses on enabling this capability of historical record-keeping. Beyond simply tracking changes, it must be able to keep track of the individual differences in the representation of the same asset in different data systems. Based on this representation, the data hub may also be used to support change propagation across systems in the future.

3.1.1. Historical Data Hub ORSD

1. **Purpose:** The purpose of the ontology is to define the key concepts used across asset management in a way that enables a history to be maintained while also enabling integration between systems such that a comprehensive asset history can be produced.
2. **Scope:** The ontology must capture assets and roles, and the ways in which they change over time. It also needs to represent the data systems and records that describe these assets.
3. **Implementation Language:** The core ontology must be defined in OWL. There is flexibility for extension in other languages to support specific reasoning tasks.
4. **Intended End-Users:**
   Asset management manager who needs access to a reliable asset history in order to support the team they oversee.

Asset management analyst who needs to know the replacement history for the various roles, and thereby the age of each asset. In addition, they need to know the history of failures, test and inspection outcomes, and repair work of the assets to infer their condition and expected useful life.

5. **Intended Uses:**
   Track changes. Track how information about assets and roles change over time, and what causes and informs the changes in a given system.
   Retrieve historical information. Retrieve information about assets and roles in the past.
6. **Non-functional requirements:**

   NFR1. The ontology must use terminology that is familiar to and easy to understand by asset management employees.

   NFR2. The ontology must enable the intuitive construction of (common) queries.

   NFR3. The task of updating data represented with the ontology (i.e. to reflect changes in knowledge) should be straightforward.

   Functional Requirements:

   CQ1. Given a change process or a defined timeframe for a given asset or role, show
   - the assets that had been involved in one of the asset change processes or
   - the roles that had been involved in one of the role change processes.

   and for each change process, show
   - which organization input the data and
   - in which data system the change was initially captured.

   CQ2. Show the status of an asset at a given time point.

   CQ3. Show the status of a role at a given time point, and if applicable, the asset serving in the role.

   CQ4. Given a point in time and an asset or role that is in a certain status, what was the change process that brought the asset or role into its status?

   CQ5. Given an asset or role involved in a particular change process, what was the status of the asset or role prior to the change process, and what was its status after the change process?

   CQ6. Given an asset status change or role status change captured on an asset or role record in a certain data system, which other data systems do not yet reflect the change?

   CQ7. Given an asset or role record in a data system, and the set of all status changes in a given timeframe, which subset of changes were made to correct mistakes in the record, versus to reflect the actual status changes?

   CQ8. Given an asset or role, what is/was known about it at a specific point in time, across all data systems?
7. **Pre-Glossary of Terms from CQs:**

- change process
- timeframe
- asset
- role
- involved in
- asset change process
- role change process
- organization
- input
- data system
- captured
- system
- timepoint
- status
- serving
- brought into
- prior to
- after
- record
- status change
- correct mistakes
- reflect changes

3.2. Content Metadata Initiative

Asset management activities are data intensive. In order for employees to do their job effectively, they need to be able to easily locate documents related to a particular asset or role. There is also a range of document types such as drawings and maintenance instructions - each may be relevant for different types of tasks. There is currently no effective way to search through the documents in the content management system at Toronto Water. Furthermore, documents that specify information about a role and an asset, such as performance requirements, are only associated with a single object (usually the asset). This means that when an asset is moved or replaced, many documents that are relevant to the role may be lost. This is a major challenge in keeping content up-to-date, as assets are constantly changing in the systems in Toronto Water's domain.

The Content Metadata Initiative focuses on using the vocabulary specified by the ontology to consistently and effectively describe stored documents. The specification of more thorough descriptions in a well-defined, agreed-upon vocabulary will enable more efficient search. One goal in particular is to develop an application that will leverage the data hub to support easy search and navigation through the content management system. The standardized vocabulary will also lend itself to metadata maintenance. Another purpose of the application will be to enable either manual or automatic update of document metadata according to changes that are recorded in other data systems. For example, if an existing asset is moved into a new role, then many of the documents related to that asset should also now be associated with the new role.

3.2.1. Content Metadata ORSD

1. **Purpose:** The purpose of the ontology is to describe the documents that are used by the Asset Management group at Toronto Water and identify the relationship between a given document and the objects (e.g. assets or roles) that it is related to to, such that relevant documents may be more easily identified by Asset Management employees.
2. **Scope:** The ontology must capture assets, roles, and documents.
3. **Implementation Language:** The core ontology must be defined in OWL. There is flexibility for extension in other languages to support specific reasoning tasks.
4. **Intended End-Users:**
   Skilled trades, analysts and reliability engineers
   who need access to documented content for specific assets, such as drawings and technical manuals, in a timely manner (under a minute). It's important that the documents are kept up-to-date, and that they don't need to spend time sifting through unrelated documents.
5. **Intended Uses:**
   Document metadata specfication. Define metadata for stored documents using the vocabulary provided by the ontology.
   Document search. Identify relevant documents through search, using the vocabulary provided by the ontology.

3.3. Data Entry Initiative

In addition to management of documents, Toronto Water employees must manage the collection and validation of data that is generated as a result of work performed – specifically, maintenance and capital project activities. This task is critical as it is the source of much of the information about an asset's life cycle. For example, these data collection efforts would describe when an asset was purchased and installed (likely part of some capital project), when it may have been repaired, and when it may have been moved (perhaps into storage, or to fulfil a different role in some other system).

The purpose of the Data Entry Initiative is to reference the specifications in the ontology to support the collection of data. The definitions specified by the ontology will be used to identify contradictions in the different systems. These contradictions may be the result of mistakes in data entry, or simply record updates that have yet to be propagated to all systems. Regardless, it is important that any contradictions may be highlighted for manual review and resolution.

Similarly, the ontology will also be used to prevent poor quality information from entering the integrated records at all by employing the same method of consistency checking *while* the user is entering the information. Feedback on the validity of a user's entry, along with reasons why it might be invalid, can be given to the user.

Finally, based on the context (e.g. what asset is being described) the ontology will be referenced to limit the data entry options available to the user. Reducing the options to the smallest possible set not only reduces opportunities for error, it simplifies the data entry task by removing the confusion, frustration, and fatigue that the irrelevant options cause in the user.

3.3.1. Data Entry ORSD

1. **Purpose:** The purpose of the ontology is to streamline the data collection process, while enabling data quality assurance and control measures.
2. **Scope:** The ontology must capture assets, roles, captial projects, as well as regular maintenance activities.
3. **Implementation Language:** The core ontology must be defined in OWL. There is flexibility for extension in other languages to support specific reasoning tasks.
4. **Intended End-Users:**
   Skilled trades who enter information specifically associated with a category of work order, such as asset condition, asset movements, and/or current operational state.
   Engineering design consultant who enters information associated with new or modified roles, such as functional capability requirements of the role.
   Construction contractor who enters information associated with new assets, including functional capability rating, and equipment classification.
5. **Intended Uses:**
   Evaluate existing data. Apply rules in the ontology to assess existing data, possibly identifying errors in the form of inconsistencies.
   Evaluate incoming data. Assess data being entered or submitted by a user to determine (and advise) whether there are any inconsistencies with respect to the ontology.
   Facilitate data entry. Reference ontology definitions to suggest input values or restrict entry options based on context (e.g. based on the asset type being described).

4. Discussion

Based on the ORSD for the Historical Data Hub Initiative alone, the following key topics have been identified as required within the asset management ontology:

Physical Assets A definition of the physical assets is required as they are central to each CQ. In particular, the CQs indicate a need to describe assets' assignment to roles, status, and involvement in processes. Different types of states are associated with assets in different data systems. At Toronto Water, the term state typically refers to a life cycle state such as "in operation" or "decommissioned". In contrast, descriptions of an asset's condition might also be interpreted as a more fine-grained state, such as "failed" or "fit for function", that is of interest when monitoring the operation of an asset.

Roles Similar to above, the CQs indicate a need to describe roles, their status and involvement in processes, and the assignment of assets to roles. In the data systems at Toronto Water, the concept of a role is closely tied, in some cases conflated, with that of an asset. Distinct from the function(s) of an asset, there is a need to represent its role in the context of the larger system. Roles in this sense are influenced by the goals of the overall system, and are affected by the elements of the system (e.g. available inputs and outputs). In contrast to the status of an asset, the status of a role may be something as simple as: "defined", "instituted", or "eliminated".

Time A number of CQs involve searching for information based on some point in time, or given “timeframe". For example, CQ2 asks for the status of an asset *at a given time point*; CQ5 asks for status information *before* and *after* some processes; and CQ7 asks for information about changes *within a given time frame*. Some representation of time and the relationships between temporal objects must be provided in order to express this aspect of the queries (and the data).

Process There is a need to capture the concept of a process and its influence on objects (e.g. assets and roles). At Toronto Water, the processes of concern range from capital projects and maintenance work, which affect assets and roles, to data reconciliation, which affects stored knowledge about assets.

Change A key requirement of this initiative is the need to capture an asset’s history - the concept of change (e.g. to an asset’s status) is critical to this.

Provenance The CQs also indicate a need to capture the origin of changes or information about changes. CQ1 asks about the *source* (creator) of information about some asset or role. Since this data might be propagated to other systems, it also asks which data system the information *originated* in. In addition, CQ7 asks about the cause or *reason* for some changes to the data.

In addition, the preliminary ORSDs indicate that the following concepts will also be required:

Information Artifacts To capture different types of documents, the structure of their contents, and relationship between the contents and individual assets or roles.

Water and Sanitation Asset Classification Classes of assets must be defined in order to support the desired data evaluation and data entry functionality.

Based on this initial specification of requirements, it appears as though a great deal of the required ontology and its intended implementation will be widely applicable for the domain of asset management. What will be distinct for Toronto Water will be the lower-level domain information: the individuals and domain-specific subclasses (e.g., specific types of water treatment assets). We view this as a promising indication that the work done on this project may be useful to support asset management groups in a variety of domains and industries.

4.1. Next Steps

There appears to be a rich body of ontology work related to asset management. However, as this work varies in its scope and purpose it is not clear whether or to what degree it may be used for this project.The next step in this work will be to complete a more detailed survey of related ontologies to identify which representations may wholly or partly satisfy the requirements for an asset management ontology. An initial speculation is that it may prove difficult to adopt ontologies that have been designed with a special application in mind, whereas general-purpose, reference ontologies such as those under development by IOF should be easier to reuse. Nevertheless, given the broad scope of asset management, it is likely that no existing work will be perfectly suited to our requirements, and so some merging, revision, and development from scratch will likely be required.

Upon completing the initial set of ORSDs, we recognize that there will be a challenge in both the specification and assessment of non-functional requirements for the ontology. These requirements may be satisfied with the design of the ontology (including structure and choice of terminology), but may also be addressed with the software in which it is implemented. The means of evaluation could range from a simple review and approval by ASMP staff, to more elaborate user testing. How to best define and evaluate these non-functional requirements (in particular, intuitiveness and ease of the query construction and data update processes) remains an open question to be considered in future work.

Further in the future, we plan to extend the data hub design to support more complex use cases. A long term goal is to develop a decision support initiative where historical information about assets and roles may be used to assist staff in determining the likely causes of some condition or symptom, identifying patterns in failures, and investigating consequences of failures.

References

[1] BSI PAS55:2008-1, Specification for the optimized management of physical assets, Standard, British Standards Institution, London, UK, 2008.

[2] The IAM competences framework, Technical Report, The Institute of Asset Management (IAM), London, UK, 2012.

[3] Framework for asset management, Technical Report, Asset Management Council Management Body of Knowledge (AMBoK), Victoria, AU, 2014.

[4] ISO 55000:2014, Asset management — Overview, principles and terminology, Standard, International Organization for Standardization, Geneva, CH, 2014.

[5] ISO 5500:2014, Asset management — Management systems — Requirements, Standard, International Organization for Standardization, Geneva, CH, 2014.

[6] J. Zeb, T. Froese, Tangible capital asset ontology in infrastructure management, Infrastructure Asset Management 1 (2014) 81–92.

[7] M. H. Karray, F. Ameri, M. Hodkiewicz, T. Louge, Romain: Towards a bfo compliant reference ontology for industrial maintenance, Applied Ontology 14 (2019) 155–177.

[8] B. S. Kulvatunyou, E. Wallace, D. Kiritsis, B. Smith, C. Will, The industrial ontologies foundry proof-of-concept project, in: IFIP International Conference on Advances in Production Management Systems, Springer, 2018, pp. 402–409.

[9] M. Hodkiewicza, E. Lowa, C. Woodsa, F. Amerib, Towards a reference ontology for maintenance work management, in: Proceedings of the Workshops of I-ESA 2020, 2020.

[10] M. C. Suárez-Figueroa, A. Gómez-Pérez, B. Villazón-Terrazas, How to write and use the ontology requirements specification document, in: OTM Confederated International Conferences" On the Move to Meaningful Internet Systems", Springer, 2009, pp. 966–982.

[11] M. Grüninger, M. S. Fox, The role of competency questions in enterprise engineering, in: Benchmarking—Theory and practice, Springer, 1995, pp. 22–31.