Description Logics for Documentation

Carlo Meghini
Consiglio Nazionale delle Ricerche
Istituto di Scienza e Tecnologia della Informazione, Pisa, Italy
carlo.meghini@isti.cnr.it

ABSTRACT

Much of the activity in a digital library revolves around collecting, organizing and publishing knowledge about the resources of the library, in the form of, *e.g.*, metadata records. In order to document such activity, digital librarians need to express knowledge about the metadata records they produce. This knowledge, which we call *documentation knowledge*, may express *e.g.*, provenance, trustability, or access restrictions of the records. Today, documentation knowledge is mostly represented in digital libraries via RDF. We propose a new type of information system, called documentation system, as a basic component of a digital library allowing to represent and reason about both domain and documentation knowledge in an expressive language such as OWL. We also show how a document system can be used in a practical setting, so as to minimize the overhead for librarians and to avoid unnecessary computational complexity.

Categories and Subject Descriptors

[**Information Systems**]: Digital libraries and archives

General Terms

Description Logic

Keywords

Documentation, Description, Metadata

1. INTRODUCTION

In a digital library, it is often necessary to represent and reason about two different kinds of knowledge. One kind is *domain* knowledge, which in is typically embodied in metadata and ontologies, used by the end-users, for instance to discover and access the resources of the library. The other kind of knowledge concerns domain knowledge and is used by digital librarians in order to manage the resources of the library. For instance, a digital librarian might want to describe the provenance, or the degree of trust or the access policy of a metadata record. We call this latter kind of knowledge as *documentation* knowledge, since documentation is one of the main reasons that brings it into life. Documentation knowledge is a primary matter dealt with by curators in libraries and archives. In general, it shows up in any organization that devotes resources to the documentation of artifacts, events, users, services and in general any resource that is of value to the organization. Its scope of application is therefore quite ample.

Documentation knowledge consists of factual and ontological statements about individuals, concepts and relations of the domain of dicourse, and as such it can be expressed and reasoned upon using standard logics, such as OWL [12]. However, a problem arises if documentation knowledge is to be used *together with* domain knowledge, as it happens in digital libraries. The problem is due to the fact that the individuals of documentation knowlegde are domain knowledge statements, and in order to express knowledge about resources and about the statements used to describe such resources *in the same language,* one needs very powerful languages, whose expressive power goes beyond that of first-order logic. Such languages, though, are hardly usable in digital libraries, because their negative computational properties are inadequate to digital library requirements.

The Resource Description Frameowrk (RDF) [9] is a case in point. RDF allows to express metadata records as statements having the described resource as subject and metadata elements as properties. It also allows to express a certain amount of documentation knowledge, by allowing properties as subjects in statements. However, as soon as the expressivity of the language goes beyond that of RDF Schema [4] by including constructs from Description Logics, serious computational problems arise. Indeed, the combination of RDF with Description Logics yields the language OWL Full [**?**], which is undecidable in spite of the decidability of its two constituents.

In this study, we tackle the problem of representing and reasoning about both domain and documentation knowlegde by means of languages significantly more expressive than RDF, but we follow a different approach than RDF. Our approach is based on the *simultaneous* usage of two different logics: the *object* logic devoted to represent domain knowledge, and the *documentation* logic (or *doc*-logic for short) devoted to represent documentation knwoledge. We view the resulting information system, which we call *documentation system,* as the backbone of a digital library, used for representing and

reasoning about domain and documentation knowledge both in an independent and in a joint way, the latter option offering an innovative query functionality. As such, the present study extends the model presented in [10], which also allows to represent documentation knowledge, but whose expressive power does not go beyond that of RDF.

In order to illustrate our approach in general terms, we choose the Description Logic $\mathcal{ALCO}$ [1] as the object logic, relying on the abstract syntax of DLs, which is more concise than the official OWL notation [?]. Section 2 presents in a more articulated way the motivations that underlay our work. Section 3 introduces the doc-logic via an example and Section 4 gives the axioms capturing part of the meaning of the concepts and roles of the doc-logic. Section 5 completes the specification of the semantics of the doc-logic, and introduces documentation systems and their querying facility. For reason of space, the types of queries that can be answered by a documentation systems are shown only via examples. Section 9 concludes by discussing the practical use of documentation systems, including computational considerations.

2. MOTIVATION
Much of the activity of digital librarians revolves around *metadata records.* Metadata records are created and maintained for many purposes, amongst which *documentation* plays a prominent role. Documentation is the process of collecting and managing evidence about a certain artifact (such as an e-book) or a certain activity (such as the lending of an e-book to a particular user) of the digital library.

A metadata record can be viewed as consisting of three separate parts: a description; a described object, the *subject* of the metadata record; and the attribution of the description to the subject.

A description is a set of features. For instance, a description in a metadata record may consist of two features: (1) being a book and (2) being titled "Waverley". Each one of these features may be understood as describing a set of resources. In particular, the former feature describes anything that is a book, while the latter describes anything that is titled "Waverley". Taken together, the two features form a description that describes any resource satisfying *both* of them. In Description Logics (DLs for short), descriptions are represented by concepts, and in fact the two terms "description" and "concept" are synonymous; we will follow the same convention. In the DL $\mathcal{ALCO}$, the above description may be represented by the concept

$$\mathsf{Book} \sqcap \exists \mathsf{Title}.\{\mathit{Waverley}\}$$

The attribution of a description to a particular subject is the act of asserting that that subject does indeed possess *simultaneously* all the features that make up the description. This attribution produces a metadata record. In a DL, attribution is realized by concept assertion, therefore in $\mathcal{ALCO}$, the metadata record as a whole can be expressed by the following concept assertion:

$$b : (\mathsf{Book} \sqcap \exists \mathsf{Title}.\{\mathit{Waverley}\})$$

where b is the DL individual that stands for the subject, a

book in this case. In a DL knowledge base (KB for short), such an assertion is placed in the ABox, the component of the KB holding factual knowledge. The axioms of DLs can be used for modelling ontologies, understood as vocabularies that establish, via terminological axioms, the meaning of the concepts used in descriptions [6]. A terminological axiom in our example, could be the following:

$$\mathit{Book} \sqsubseteq (\exists \mathsf{Title.Literal} \sqcap \exists \mathsf{Author.Person})$$

stating that every book has a title that is a literal, and an author that is a person. The ontology underlying an ABox is placed in the other component of a DL KB, the TBox.

The discussion so far seems to indicate that DLs are adequate representation languages for the knowledge contained in a digital library. Indeed, OWL [12], the W3C recommendations that stand for DLs in the semantic web architecture [3], is known and used in digital libraries.

However, there is a fundamental requirement of digital libraries that DLs are not able to capture. This requirement comes from the fact that documentation, a central activity in a digital library, may be a *recursive* process; it may not be solely confined to the domain of discourse, but it may also concern the descriptions and metadata records used for documenting the resources of the domain of discourse. For example, a digital librarian may need to represent the fact that a metadata record has a certain provenance, (*i.e.,* it has been created by a certain person, on a certain date, as a result of a certain activity), or a certain degree of trustability; or that the record is subject to a certain set of access restrictions, or to a certain billing policy.

These requirements are not exclusive of digital librarians. Other curators, in archives, museums and galleries, deal with possibly different kinds of objects than those found in digital libraries, yet have the same requirements concerning the documentation of their objects or of the activities that involve them. In general, the same requirements arise in every information system that deals with the documentation of the individuals of the domain of discourse, be these individuals users, devices, or events.

In order to cope with the recursiveness of documentation, we argue that a knowledge representation language is needed that offers descriptions as first class citizens, that is as individuals on their own right, endowed with an identity and a structure. In addition, the language in question should offer the machinery to attribute a description to a certain individual, which may itself be a description, thereby coping with the recursiveness of documentation.

DLs, in spite of their name, fall short of this requirement. They offer a rich machinery to create descriptions, yet these descriptions are not denotable as individuals, and therefore it is not possible to state any knowledge on them, other than assertions and axioms. As such, the support that DLs offer to documentation is practically limited.

The situation, however, is not unrecoverable. The machinery of DLs can be used to remedy at this inconvenient, by slightly shifting the focus of representation from the individuals, concepts and roles in the domain of discourse, to

descriptions made up of these. The shift is a form of reification and leads to a logic, the doc-logic, that allows the expression of documentation knowledge, in contrast to the object-logic that is used for representing domain knowledge.

In order to show how this can be done, in the next Section we introduce a doc-logic corresponding to the object-logic $\mathcal{ALCO}$. Our doc-logic will be doc-$\mathcal{ALCO}$, more simply called as *alco*.

3. INTRODUCING $\mathcal{ALCO}$

Suppose we wish to create the above description for the Waverley, and we also want such description to be identified as the individual d.[1] In order to achieve our goal, in *alco* we describe the structure of d as the conjunction of two concepts, which we choose to identify as d_1 and d_2. We use the *alco* role CCId (for *Concept Conjunction Identification*) for associating the identifier of a conjunction to the identifiers of the conjuncts, as follows:

$$\mathsf{CCId}(d, d_1) \quad \mathsf{CCId}(d, d_2)$$

The role CCId binds an identifier to a *part* of a concept, and therefore we will call it a *binding* role.

Next, we need to state that d_1 identifies an atomic concept, say *Book*. For this, we use the *alco* role ACId (for *Atomic Concept Identification*) as follows:

$$\mathsf{ACId}(d_1, Book)$$

An assertion on the role ACId is in fact an assigment of an identifier to a whole concept, not just to a part of it, as in the case of binding roles; therefore will call ACId a *concept identifying role,* or more simply an *identifying role.* Of course, we could choose to use an atomic concept as an identifier of itself (*e.g.*, the *alco* individual *Book* as an identifier of the $\mathcal{ALCO}$ concept *Book*) to make a doc-KB more readable; but we prefer to use different names in order to avoid confusion. We also introduce in *alco* the role CNId (for *Concept Negation Identification*) for identifying negation concepts as follows: $\mathsf{CNId}(d, e)$ assigns the identifier d to the negation of the concept identified by e. Also CNId is an identifying role.

As seen in the previous section, the title feature is expressed in $\mathcal{ALCO}$ as the concept $\exists Title.\{Waverley\}$, known in the DL world as *existential quantification.* In order to reduce this concept to something denotable as the individual d_2, we follow the same style followed above for concept conjunction, as follows:

- we introduce the identifying role ARId (for *Atomic Role Identification*), and use it to identify the $\mathcal{ALCO}$ role *Title* as d_3 by the assertion $\mathsf{ARId}(d_3, Title)$;

- we introduce the identifying role SCId (for *Singleton Concept Identification*) and use it to identify the $\mathcal{ALCO}$ singleton concept $\{Waverley\}$ as d_4 by the assertion $\mathsf{SCId}(d_4, Waverley)$;

---

[1] We are aware that in real digital libraries, descriptions are created by using terms from formal vocabularies, such as Dublin Core; however, this aspect is entirely irrelevant to our present purposes, therefore we opt for a better readability of our examples by using terms from the English vocabulary.

- finally, we introduce two identifying roles ERId (for *Role Identification in an Existential Quantification*) and ECId (for *Concept Identification in an Existential Quantification*) and use them to bind d_2 to its constituent parts by the assertions:

$$\mathsf{ERId}(d_2, d_3) \quad \mathsf{ECId}(d_2, d_4)$$

We remark that we could use a ternary role for modeling the existential quantification; as it is well known, DLs with n-ary roles are generalizations of DLs, and the generalization does not impact on the complexity of reasoning in any significant way. However, we prefer to stay with the more familiar basic machinery of DL in order to fully highlight the features of *alco,* also in terms of axioms (see next section).

Now we have completed the description of d, and can create the desired metadata record by attributing d to the book b that we want to describe. The obvious way to state the association between d and b, is to introduce the *alco* role CAss (for *Concept Assertion*) and use it as follows:

$$\mathsf{CAss}(d, b)$$

The first argument of a CAss assertion is always a description identifier, whereas the second argument identifies the subject, which can be any resource. The CAss role allows to represent in *alco* an $\mathcal{ALCO}$ concept assertion. Yet, it is not adequate to document metadata record, because there is no individual that denotes the resulting metadata record in a CAss assertion. An alternative way of proceeding is to coin a name, say m, for the metadata record that we want to create, and then connect m to d and b by using appropriate role assertions, namely:

$$\mathsf{MRD}(m, d) \quad \mathsf{MRS}(m, b)$$

There is an obvious relation between the just introduced roles and CAss. This relation will be captured by an axiom, introduced in next Section.

So far we have simply used a rather cumbersome notation for expressing a concept assertion. However, this notation has given us identifiers for the description and the metadata record that we have created, therefore we are in the position of representing documentation knowledge about them. Suppose we want to state that d was created by *John.* The latter feature can be expressed as a description identified by e as follows:

$$\mathsf{ARId}(e_1, Author) \quad \mathsf{SCId}(e_2, John) \quad \mathsf{ERId}(e, e_1) \quad \mathsf{ECId}(e, e_2)$$

The description e can be attributed to d by stating

$$\mathsf{CAss}(e, d)$$

Analogously we can document the metadata record m by specifying a creation time, a creation place and an author for it. We first create a description h for the $\mathcal{ALCO}$ concept $\exists Author.\{Sue\}$, using the *alco* roles (for brevity, we do not detail h); finally, we add the assertion:

$$\mathsf{CAss}(h, m)$$

to the ABox of the doc-KB.

Role	Meaning
ACId(d, A)	d identifies the atomic concept A
ARId(d, R)	d identifies the atomic role R
CCId(d, e)	the concept identified by e is a conjunct of the concept identified by d
CNId(d, e)	d identifies the negation of the concept identified by e
ERId(d, R)	d identifies an existential quantification having R as role
ECId(d, e)	d identifies an existential quantification whose concept is identified by e
SCId(d, o)	d identifies the singleton concept $\{o\}$
CAss(d, e)	the individual e is an instance of the concept identified by d
CSAx(d, e)	the concept identified by d is subsumed by the concept identified by e
MRD(m, d)	d is the description of the metadata record m
MRS(m, b)	b is the subject of the metadata record m

Table 1: The *alco* roles

In order to make *alco* a full doc-logic, we need a role for representing (at the doc-level) concept subsumptions in an object-TBox. To this end, we introduce CSAx (for *Concept Subsumption Axiom*) with the intended meaning that:

$$\text{CSAx}(e, d)$$

states that the concept identified by e is a sub-concept of the concept identified by d.

Summing up the expressive machinery of *alco,* we can divide *alco* roles in the following categories (Table 1 recaps the *alco* roles and their informal meaning):

- roles for representing *descriptions,* that is concepts of the object-DL; these roles have been already classified into binding and identifying roles and reflect the concept constructors of the object-DL, which is $\mathcal{ALCO}$;

- the CAss role for representing concept assertions of the object-DL; we note that concept assertions can also be used to express role assertions, since $a : \exists R.\{b\}$ is equivalent to $R(a, b)$. For this reason there is no machinery in *alco* to represent role assertions of the object-DL;

- the CSAx role for representing inclusion axioms of the object-DL;

- roles for representing metadata records; these are the MRD and the MRS roles.

As it can be seen, the first three categories above are the doc-machinery proper, because the roles that they include are used to represent the components of an object KB at the doc-level. The last category exploits the doc-machinery to the end of representing metadata records, that are the core notion of documentation.

We note that the concept constructors modeled by *alco* include unary (*e.g.*, negation or singleton), binary (*e.g.*, con-

junction) and ternary (*e.g.*, existential quantification) operators, therefore *alco* exemplifies the machinery required to deal at the doc-level with *any* DL, even though different concept constructors would require different doc-logic roles.

We conclude this section by indicating other kinds of knowledge that are relevant to the activity of documentation, and that can be represented in *alco.* In particular, in *alco* it is possible to:

- Represent lexical knowledge about the terms used for concepts or roles in descriptions (such as *Author* or *Book*), as resources having one or more lexical forms (such as "Author" or "Autore"), and related to other resources in the same lexical space by relations such as "similar to" or related to", typically found in *thesauri.* This knowledge is useful for facilitating the users of the digital library who use terms to designate resources that are not completely aligned with those used in descriptions. This knowledge belongs to the lexicon of the domain of discourse along with ontological knowledge, but, unlike ontological knowledge it cannot always be expressed via DL axioms because it has a different logical nature [7]. Today, lexical knowledge is expressed in digital libraries by using on the SKOS vocabulary [?], but in order to do so properties and classes have to be seen as objects (technically, as instances of the SKOS class Concept) thereby introducing ambiguity in the representation.

- Create collections of metadata records which are at the basis of *catalogs.*

- Represent knowledge about the actions that concern descriptions and metadata records. These actions are special kinds of events, performed by digital librarians in their daily activity, and their expression would allow to document such activity in an appropriate way.

For reasons of space, we no longer discuss these usages of *alco.*

4. SEMANTICS
As a DL, *alco* has a well-defined semantics, based on the notion of interpretation. This notion, however, turns out to license undesired situations. In this section we discuss these undesired situations, and introduce axioms for ruling them out. These axioms are intended to be in the TBox of *any* doc-KB, since they express the semantics of the *alco* roles. In order to ease the expression of the axioms, we introduce some atomic concepts:

- OInd, OACon and OARol, for representing individuals, atomic concepts and atomic roles of the object-DL, respectively; these concepts are defined as follows (as customary, $\exists R$ abbreviates $\exists R.\top$):

$$\text{OInd} \equiv \exists \text{SCId}^-; \quad \text{OACon} \equiv \exists \text{ACId}^-; \quad \text{OARol} \equiv \exists \text{ARId}^-$$

- Concepts denoting the different types of concept and role identifiers:

$$\text{AtomId} \equiv \exists \text{ACId}; \quad \text{ConjId} \equiv \exists \text{CCId}; \quad \text{NegId} \equiv \exists \text{CNId}$$
$$\text{SomeId} \equiv \exists \text{ERId}; \quad \text{SingId} \equiv \exists \text{SCId}; \quad \text{RoleId} \equiv \exists \text{ARId}$$

- The concept ConcId denoting concept identifiers:

$$\mathsf{ConcId} \equiv \mathsf{AtomId} \sqcup \mathsf{ConjId} \sqcup \mathsf{NegId} \sqcup \mathsf{SomeId} \sqcup \mathsf{SingId}$$

4.1 Identification axioms

The first group of axioms capture the proper behaviour of the binding and the identifying roles.

First, identifying roles must behave like functions (in OWL terms, they are functional object properties [11]).

$$(\geq 2\ \mathsf{ACId}) \sqsubseteq \bot; \quad (\geq 2\ \mathsf{ARId}) \sqsubseteq \bot; \quad (\geq 2\ \mathsf{CNId}) \sqsubseteq \bot$$
$$(\geq 2\ \mathsf{ERId}) \sqsubseteq \bot; \quad (\geq 2\ \mathsf{ECId}) \sqsubseteq \bot; \quad (\geq 2\ \mathsf{SCId}) \sqsubseteq \bot$$

Concerning the binding role CCId, every identifier used as a first argument in an assertion on this role, must appear as a first argument also in at least another assertion on the same role, because a conjunction has at least two conjuncts. We express this condition as follows:

$$(= 1\ \mathsf{CCId}) \quad \sqsubseteq \quad \bot$$

Moreover, for every $\mathsf{ERId}(d, e)$ assertion, there must be exactly one $\mathsf{ECId}(d, e')$ assertion, for any individuals e, e'; and viceversa. In order to capture this constraint, which we call *pairing constraint,* it is not sufficient to include the axiom $\exists\mathsf{ERId} \equiv \exists\mathsf{ECId}$ in the TBox. This is due to the fact that a DL KB is interpreted under the Open World Assumption. As a consequence, a KB whose ABox contains *only* the assertion $\mathsf{ERId}(d, R)$ and whose TBox contains the axiom $\exists\mathsf{ERId} \equiv \exists\mathsf{ECId}$ is not inconsistent; rather, the KB is understood as implicitly stating that there exists some unknown individual z such that $\mathsf{ECId}(d, z)$ is true in every interpretation. We cannot therefore capture the pairing constraint as a TBox axiom; we will do it in a different way, illustrated in the next Section.

Finally, binding and identifying roles must all together satisfy the obvious constraint that the domain of each one of them be disjoint from the domain of each of the others, otherwise it may happen that the same identifier be used to identify two concepts of different kinds. Assuming the above mentioned pairing constraints are in place, we need to consider only one of ERId and ECId because the domains of these two roles are made equivalent by the pairing constraint. So, overall we need to declare mutual disjointness of the domain of six roles; this requires fifteen axioms, all of the same type. For brevity, we only state the five axioms stating the disjointness of atomic concept identifiers from the other types of concept identifiers:

$$(\mathsf{AtomId} \sqcap \mathsf{ConjId}) \sqsubseteq \bot; \quad (\mathsf{AtomId} \sqcap \mathsf{NegId}) \sqsubseteq \bot;$$
$$(\mathsf{AtomId} \sqcap \mathsf{SomeId}) \sqsubseteq \bot; \quad (\mathsf{AtomId} \sqcap \mathsf{SingId}) \sqsubseteq \bot;$$
$$(\mathsf{AtomId} \sqcap \mathsf{RoleId}) \sqsubseteq \bot$$

4.2 Syntactic axioms

Syntactic axioms are those making sure that the syntax of the object-DL concepts is properly captured by the assertions of the doc-DL.

A basic constraint of every DL, is that the symbols used for individuals, atomic concepts and roles come from three disjoint alphabets.

$$(\mathsf{OACon} \sqcap \mathsf{OARol}) \sqsubseteq \bot; \quad (\mathsf{OACon} \sqcap \mathsf{OInd}) \sqsubseteq \bot;$$
$$(\mathsf{OARol} \sqcap \mathsf{OInd}) \sqsubseteq \bot$$

We also must make sure that concepts are properly formed in a doc-KB. To exemplify, if in the doc-KB there is the assertion that d identifies a negation, *i.e.*, $\mathsf{CNId}(d, d_1)$, then d_1 must identify something, and a concept in particular. The second part of this constraint can be captured via an axiom, imposing that anything that appears as a second argument in a CNId assertion be a concept identifier, and the same for all the other concept constructors. This can be done by introducing the following axioms:

$$\top \sqsubseteq (\forall\mathsf{CCId}.\mathsf{ConcId}); \quad \top \sqsubseteq (\forall\mathsf{CNId}.\mathsf{ConcId});$$
$$\top \sqsubseteq (\forall\mathsf{ECId}.\mathsf{ConcId})$$

However, the first part of the constraint, namely that there be in the KB an explicit assertion binding d_1 to a concept, cannot be formalized as an axiom, again due to the Open World Assumption. Similarly to pairing constraints, we have therefore to find a different way of expressing this kind of constraints, which we call the *syntactic constraints.*

4.3 Metadata axioms

Metadata axioms concern identifiers of metadata records. Our intended notion of a metadata record requires that a metadata record concern exactly one individual. In order to capture this intention, we introduce the following axiom:

$$(\geq 2\ \mathsf{MRS}) \quad \sqsubseteq \quad \bot$$

Moreover, we need to enforce a pairing constraint for MRD and MRS, in the sense that for every assertion of the kind $\mathsf{MRD}(m, d)$, we want the KB to contain at least one assertion of the form $\mathsf{MRS}(m, i)$, for some individual i; and viceversa. For the reasons given above, also this kind of pairing constraint has to be captured in a different way.

Finally, we include an axiom for capturing the previously mentioned relation between the role MRD and MRS from one side, and the role CAss from the other side. In fact, every time we use MRD and MRS for structuring a metadata record m, as in $\mathsf{MRD}(m, d)$ and $\mathsf{MRS}(m, b)$, we are implicitly asserting that the d describes b, that is, that $\mathsf{CAss}(d, b)$. In order to make this connection happen in a doc-KB, we introduce the axiom:

$$\mathsf{MRD}^- \circ \mathsf{MRS} \quad \sqsubseteq \quad \mathsf{CAss} \tag{1}$$

where $\circ$ is the role composition operator. Notice that this axiom leaves the freedom of inserting CAss assertions without the corresponding MRD and MRS assertions. In other words, it is possible to create concept assertions that are not metadata records, such as for instance the last assertion of the previous example.

4.4 Inference Axioms

Inference axioms are required in order to model the proper behavior of the CAss and the CSAx roles.

Concerning the CAss role and returning to our example, from the assertions $\mathsf{CAss}(d, b)$ and $\mathsf{CCId}(d, d_1)$ it should follow the assertion $\mathsf{CAss}(d_1, b)$, because d identifies a concept

conjunction and d_1 identifies one of the conjuncts. By the same argument, it also should follow $\mathsf{CCId}(d, d_2)$. More generally, modelling the semantics of CAss means laying down all the rules that capture implied concept assertions, ultimately leading to the axiomatization of (object-level) instance checking in the doc-logic.

Likewise, modelling the semantics of the CSAx role amounts to axiomatize (object-level) concept subsumption in the doc-logic. In other words, the doc-logic does not only have to capture the syntax of the corresponding object-logic, but also its inference mechanism.

The proof theory of DLs provides us with sound and complete inference methods for both instance checking and subsumption. These methods can be encoded in *alco* by means of axioms and semantic conditions, exactly in the same way the syntax rules of the object-logic are encoded. As a result, the users of the doc-KB would be able to exploit the implicit domain knowledge, typically for performing ask operations. In particular, in order to check whether an individual i is an instance of an object-concept c, it suffices to identify c via an individual d and then ask whether $\mathsf{CAss}(d, i)$ logically follows from the doc-KB. The same machinery also allows to do some consistency checking, for instance checking whether an individual is an instance of the $\perp$ concept, or of a contradictory description. However, the price to be paid for achieving this goal would be very high, as the encoding of inference would make the doc-logic very complex.

In what follows, we will present a much simpler method for attaining the same goal. Our method is based on the separation of concerns: implicit documentation knowledge is to be extracted from the doc-KB, wehreas implicit domain knowledge is to be extracted from the object-KB, and the same applies to consistency checking. A detailed account of our method will be presented below. Here we just note that this approach dispenses us from modelling object-level inference in the doc-logic, keeping the latter at a much simpler level.

In order to preserve the semantics of the *alco* roles, we assume that the doc TBox does not contain any axiom concerning these roles other than those given above.

5. STRONG CONSISTENCY

As we have already observed, the axioms that we have introduced so far do not capture pairing and syntactic constraints. As a result, a consistent doc-KB may violate these constraints, defeating intuition. But there is also another intuitive notion of inconsistency that escapes the axioms, as illustrated by the following example. Suppose we want to add to the KB the assertion that b is not a book. In $\mathcal{ALCO}$, this would be accomplished by introducing the concept assertion:

$$b \; : \; \neg Book$$

in the ABox. In *alco,* instead, the complex concept $\neg Book$ must first be created and identified, say as the individual f, by using the *alco* roles, and then f must be declared to be the negation of the atomic concept *Book*. There is already an identifier for the latter concept, namely d_1, therefore all our replacement librarian must do, is to state:

$$\mathsf{CNId}(f, d_1)$$

so that he can finally assert that b is an instance of f, by using the *alco* role assertion:

$$\mathsf{CAss}(f, b)$$

By adding the last two assertions to those introduced in the Section 3, the consistency of the doc-KB is not broken. However, the represented knowledge is clearly inconsistent from an intuitive point of view, since b is asserted to be a book and not a book at the same time.

One way to capture pairing and syntacting constraints while at the same time ruling out the last kind of inconsistency, is to try to transform a consistent doc-KB into its corresponding object-KB. In order to do so, $\mathcal{ALCO}$ concepts must be extracted from the assertions of the doc-ABox. If the doc-KB suffers from a pairing or a syntactic inconsistency, then this extraction is not possible because the assertions in the doc-ABox do not conform to the syntax of the doc-logic. Otherwise the extraction is possible, and it allows to determine, for each concept identifier d in the doc-ABox, the concept identified by d, that we denote as $\nu_D(d)$ ($\nu(d)$ for simplicity). Once the function ν is determined, the doc-KB can be transofrmed into its corresponding object-KB, and the obtained object-KB can be checked for consistency. If this check succeeds, then the doc-KB satisfies all the intuitive consistency criteria.

This approach gives us a simple method for checking consistency of the doc-KB. A further advantage of it, is that we no longer need to model object level instance checking or subsumption as implicit CAss and CSAx assertions, respectively; we can transform the doc-KB and perform these inferences on the resulting object-KB, by relying on well-known algorithms.

This is the route that we will follow in the rest of the paper. To this end, we will first define how to determine the function ν from a given doc-KB D. Based on the existence of $\nu(d)$ for each concept identifier d, we will define a stronger consistency criterion for a doc-KB. Next, we state the transformation ϕ from a doc-KB to its corresponding object-KB and define the strongest consistency criterion of a doc-KB based on the consistency of its corresponding object-KB.

The domain of the function ν is the set of concept identifiers:

$$dom(\nu) = \{d \mid (T, A) \models d : \mathsf{ConcId}\}$$

and can be efficiently determined as the set of identifiers that occur as first arguments in a binding or in an identifying role assertion. For each $d \in dom(\nu)$, the value of $\nu(d)$ is recursively defined in Table 2. By iterating this recursive computation on $dom(\nu)$, the function ν can be efficiently computed. To exemplify, we show the derivation of the concept $\nu(d)$ associated to identifier d in our running example. We have:

1. the doc-ABox contains $\mathsf{CCId}(d, d_1)$ and $\mathsf{CCId}(d, d_2)$, therefore $\nu(d) = \nu(d_1) \sqcap \nu(d_2)$

2. the doc-ABox contains $\mathsf{ACId}(d_1, Book)$, therefore $\nu(d_1) = Book$

3. the doc-ABox contains $\mathsf{ERId}(d_2, d_3)$ and $\mathsf{ECId}(d_2, d_4)$, therefore $\nu(d_2) = \exists \nu(d_3).\nu(d_4)$

$\nu(d)$	if the doc-ABox contains
A	$\mathsf{ACId}(d, A)$
R	$\mathsf{ARId}(d, R)$
$\nu(e_1) \sqcap \ldots \sqcap \nu(e_n)$	$\mathsf{CCId}(d, e_1), \ldots, \mathsf{CCId}(d, e_n)$, n maximal
$\neg \nu(e)$	$\mathsf{CNId}(d, e)$
$\exists \nu(e).\nu(f)$	$\mathsf{ERId}(d, e)$ and $\mathsf{ECId}(d, f)$
$\{o\}$	$\mathsf{SCId}(d, o)$
ω	otherwise

Table 2: Assignment of object-DL concepts to identifiers

4. the doc-ABox contains $\mathsf{ACId}(d_3, Title)$, therefore $\nu(d_3) = Title$

5. the doc-ABox contains $\mathsf{SCId}(d_4, Waverley)$, therefore $\nu(d_4) = \{Waverley\}$

Overall, we have $\nu(d) = Book \sqcap \exists Title.\{Waverley\}$

Intuitively, if $D = (T, A)$ is a consistent doc-KB, then exactly one of the conditions on the right column of Table 2, excluding the last row, is met by the ABox A. In fact, consistency of the doc-KB guarantees that there is only one ACId assertion having d as first term, and that no other role assertion has d as first term. And the same for the other roles. Moreover, if pairing and syntactinc constraints are satisfied, then the last row of the table is never entered, because every CCId, ERId and ECId assertion is paired by a corresponding one, and every used identifier is associated with a concept. Overall, therefore, if the doc-KB is consistent and if pairing and syntactinc constraints are satisfied, we have:

$$\nu(d) \neq \omega \text{ for all } d \in dom(\nu),$$

in other words ν *is total on D*. Based on this consideration we define a doc-KB $D = (T, A)$ to be *fully consistent* if $\nu(d)$ is total on D.

Let us now consider how to define the transformation ϕ. The DL *alco* offers the roles CSAx and CAss for representing the terminological axioms and the assertions of the object-KB, respectively. Therefore, it is natural to use assertions on the former role in order to derive the axioms in the object-TBox, and assertions on the latter role in order to derive the axioms in the object-ABox. Formally, given a doc-KB $D = (T, A)$, we have:

$$\begin{aligned} \phi(D) &= (\mathcal{T}, \mathcal{A}) \\ \mathcal{T} &= \{\nu(d) \sqsubseteq \nu(e) \mid \mathsf{CSAx}(d, e) \in A\} \\ \mathcal{A} &= \{i : \nu(d) \mid (T, A) \models \mathsf{CAss}(d, i)\} \end{aligned}$$

Some explanations are in order concerning the translation of CAss assertions. Any such assertion has the form $\mathsf{CAss}(d, i)$ where d is a description identifier and i is either a description identifier (such as e in the example above) or an individual denoting any other kind of resource (such as the book b in the example above). In the latter case, i does not have to be translated, as there would be nothing to translate it into. In the former case, i is a concept identifier and d another concept identifier that describes i. If we translate both identifiers into the corresponding concepts, the result

doc-ABox	ν	object-ABox
$\mathsf{CCId}(d, d_1)$, $\mathsf{CCId}(d, d_2)$	$d \mapsto \nu(d_1) \sqcap \nu(d_2)$	
$\mathsf{ACId}(d_1, Book)$	$d_1 \mapsto Book$	
$\mathsf{ERId}(d_2, d_3)$, $\mathsf{ECId}(d_2, d_4)$	$d_2 \mapsto \exists \nu(d_3).\nu(d_4)$	
$\mathsf{ARId}(d_3, Title)$	$d_3 \mapsto Title$	
$\mathsf{SCId}(d_4, Waverley)$	$d_4 \mapsto \{Waverley\}$	
$\mathsf{CAss}(d, b)$		$b : Book \sqcap$ $\exists Title.\{Waverley\}$
$\mathsf{MRD}(m, d)$, $\mathsf{MRS}(m, b)$		
$\mathsf{ERId}(e, Author)$, $\mathsf{ECId}(e, e_1)$	$e \mapsto \exists Author.\nu(e_1)$	
$\mathsf{SCId}(e_1, John)$	$e_1 \mapsto \{John\}$	
$\mathsf{CAss}(e, d)$		$d : \exists Author.\{John\}$
$\mathsf{ERId}(h, Author)$, $\mathsf{ECId}(h, h_1)$	$h \mapsto \exists Author.\nu(h_1)$	
$\mathsf{SCId}(h_1, Sue)$	$h_1 \mapsto \{Sue\}$	
$\mathsf{CAss}(h, m)$		$m : \exists Author.\{Sue\}$

Table 3: Summary of the running example

will be an assertion like $C : D$ where both C and D are object-concepts. In our example, the translation would look like:

$$(Book \sqcap \exists Ttile.\{Waverley\}) : \exists Author.\{John\}$$

The natural language reading of the last expression is "John created the description *Book titled 'Waverley'*", which is precisely what we want to say. However, the above is not a valid assertion in any DL, and this is the very reason why we set out to define the doc-DL. Therefore, we have no choice but to translate only the second identifier. This will result in the concept assertion:

$$d : \exists Author.\{John\}$$

This means that d will appear in the object-KB, but its connection with the description that it identifies (*i.e.*, $(Book \sqcap \exists Ttile.\{Waverley\})$) is lost in the object-KB.

Notice that we require $\mathsf{CAss}(d, i)$ to be *not only* explicitly asserted, but also implicitly present in the KB, typically as a consequence of the creation of a metadata record.

Table 3 shows the translation of the doc-KB of our running example.

Based on these considerations, we say that a doc-KB $D = (T, A)$ is *strongly consistent* if D is fully consistent and $\phi(D)$ is a consistent $\mathcal{ALCO}$ KB.

6. DOCUMENTATION SYSTEMS

A *documentation system $\mathcal{S}$* is a pair $\mathcal{S} = \langle D, O \rangle$ where D is a doc-KB and O is an object-KB, such that $O = \phi(D)$.

The definition captures the intuitive notion that a documentation system (DS for short) is a two-sided information system: it has an object side, offering functionality on domain knowledge to end users; and it has a documentation side, offering functionality on documentation knowledge to curators. These two sides must be connected in the intuitive way, and the machinery developed so far aims at guaranteeing from the formal point of view that this is indeed the case.

We expect a DS to be the backbone of a larger system, storing a wide range of knowledge on both sides. On the object side, the DS may store knowledge about the entities involved

in descriptions, such as authors and subjects of books, to stay with our library example. On the documentation side, the DS may store knowledge about the resources of the organization, such as loans and bills. These additional kinds of knowledge can clearly be accommodated in a DS without altering the structure defined above.

The presence of both domain and documentation knowledge in a DS, opens up interesting scenarios in querying, as discussed in the next Section.

6.1 Querying a documentation system

The two KBs in a DS can be queried individually, based on the kind of knwoledge they store. To illustrate, let us consider our running example. An end user in need to know the different editions of the *Waverley* in the digital library, must obviously ask the object-KB, where the resources of the library are described. We will call *object-query* any query that can be answered by the object-KB. In our example, a query might be something like $Book \sqcap \exists Title.\{Waverley\}$. As customary, we will denote the answer to an object query C_o against an object-KB O as $\text{Ask}_o(C_o, O)$, or simply $\text{Ask}_o(C_o)$, and define:

$$\text{Ask}_o(C_o) = \{ \, i \mid O \models i : C_o \}$$

Returning to the example, the doc-ABox contains both $\text{MRD}(m,d)$ and $\text{MRS}(m,b)$, and therefore, due to the axiom 1 in the doc-TBox, we have that $D \models \text{CAss}(d,b)$. Based on the translation function ϕ, the assertion:

$$b : Book \sqcap \exists Title.\{Waverley\}$$

ends up in the object-ABox, therefore $b \in \text{Ask}_o(C_o)$.

On the other hand, a library manager in need to know which metadata records exist for a certain resource, must ask the doc-KB, formulating their need as a *doc-query.* In our example, a doc-query C_d might be something like $\exists \text{CAss}.\{b\}$, asking for the metadata records about b. We will use $\text{Ask}_d(C_d)$ for denoting the result of asking doc-query C_d to a doc-KB and define:

$$\text{Ask}_d(C_d) = \{ \, i \mid O \models i : C_d \}$$

Clearly, $d \in \text{Ask}_d(\exists \text{CAss}.\{b\})$.

There is a third category of queries, which we call *mixed* queries, that can be asked to a DS. Mixed queries involve both domain and documentation knowledge, and can be of one of two kinds:

- *mixed-object* queries, asking for the resources of the object-KB that satisfy some property expressed in the doc-KB; for instance, a mixed-object query may ask for the editions of the *Waverley* that have been described by John;

- *mixed-doc* queries, asking for the resources of the doc-KB that satisfy some property expressed in the object-KB; for instance, a mixed-object query may ask for the metadata records of the *Waverley* that have been created by Sue.

In order to express mixed queries, we need to address both the object- and the doc-KB. One way of doing so, is to use the operators Ask_o and Ask_d *within* queries. Using these two operators, the editions of the *Waverley* described by John can be expressed as follows:

$$Book \sqcap \exists Title.\{Waverley\} \sqcap \text{Ask}_d(\exists \text{CAss}^-.\text{Ask}_o(\exists Author.\{John\}))$$

This is a concept denoting the individuals in the object-KB that are known to be books titled *Waverley,* and that are known in the doc-KB to be described by a description authored by *John.* Notice the double nesting of the Ask operator, without which it would not be possible to express the query. This implies that asking a doc-KB may require asking a object-KB. Let us now consider the mixed-doc query asking for the metadata records of the *Waverley* authored by Sue. Using the same ask operators introduced above, this query can be expressed as:

$$\exists \text{MRD}.(\exists \text{CAss}.\text{Ask}_o(\exists Title.\{Waverley\})) \sqcap \text{Ask}_o(\exists Author.\{Sue\})$$

This is a conjunction of two concepts: the first concept denotes the metadata records whose description describes an individual known in the object-KB to be titled *Waverley.* The second concept denotes the individuals known in the object-KB to be authored by Sue. The conjunction of these concepts clearly expresses the desired query.

We can formally define the query language used in the above examples as a concept language including two types of concepts:

1. object-concepts, for expressing object-queries; object-concepts are an extension of $\mathcal{ALCO}$ concepts with the Ask_d operator introduced above;

2. doc-concepts, for expressing doc-queries; doc-concepts are a contraction of *alco* concepts, excluding numerical restrictions and disjunction, but adding the Ask_o operator introduced above.

The syntax of our concept language is given by the following rules, where the subscripts o and d are used for elements of the object- and the doc-DL, respectively.

$$
\begin{aligned}
C &::= C_o \mid C_d \\
C_o &::= \text{any object-concept} \mid \text{Ask}_d(C_d) \\
C_d &::= \text{any doc-concept} \mid \text{Ask}_o(C_o)
\end{aligned}
$$

We call this language Documentation Query Language (DQL for short). It is clear that any arbitrary set of DL constructors can be employed on either side of DQL.

In order to give the semantics of DQL, the intepretations of the doc- and the object-DL are combined in a rather obvious way. For simplicity, we omit such specification, pointing out that it allows us to reduce query answering on a DS to query answering on a single KB whose TBox is given by the union of the TBoxes of the doc- and the object-KBs, and the same for the ABox. In order to show the effects of this result, let us now return to our example queries. From a semantical point of view, the former query above: $Book \sqcap \exists Title.\{Waverley\} \sqcap \text{Ask}_d(\exists \text{CAss}^-.\text{Ask}_o(\exists Author.\{John\}))$ stated against our example DS in Table 3, is equivalent to the query $Book \sqcap \exists Title.\{Waverley\} \sqcap \exists \text{CAss}^-.\exists Author.\{John\}$

stated against the union of the KBs shown in the first and third column of the table. The union contains the assertions:

$$b : Book \sqcap \exists Title.\{Waverley\}$$
$$\mathsf{CAss}(d,b)$$
$$d : \exists Author.\{John\}$$

clearly implying that b is in the answer to the query.

Analogously, the latter query above:
$\exists MRD.(\exists \mathsf{CAss}.\mathrm{Ask}_o(\exists Title.\{Waverley\})) \sqcap \mathrm{Ask}_o(\exists Author.\{Sue\})$
stated against $\mathcal{S} = (D, O)$, is equivalent to the query
$\exists MRD.(\exists \mathsf{CAss}.\exists Title.\{Waverley\}) \sqcap \exists Author.\{Sue\}$ stated against $D \cup O$. The assertions:

$$\mathsf{MRD}(m,d) \quad \mathsf{CAss}(d,b)$$
$$b : Book \sqcap \exists Title.\{Waverley\}$$
$$m : \exists Author.\{Sue\}$$

are in $D \cup O$, therefore m is in the answer to the query.

7. USABILITY

We are now in the position of stating where *alco* is placed in the $\mathcal{AL}$ family. In order to express the doc-logic axioms, we used as concept constructors: conjunction, disjunction, universal quantification, unqualified existential quantification, unqualified numerical restrictions, $\top$ and $\bot$; as role constructors, we used inverse roles and role composition. The simplest member of the $\mathcal{AL}$ family that includes all the *alco* construcotrs is the $\mathcal{ALCQI}_{reg}$ DL, a very expressive logic, of whose expressive power we have used only a small part (for instance, we have not used full negation or transitive roles). As such, *alco* is a decidable logic (based on Theorem 5.16 in [1], logical implication in this logic is ExpTime-complete), but clearly intractable.

However, this negative computational behavior is relevant only if the expressive power of *alco* is used for constructing rich doc-KBs that exploit the full expressive power of $\mathcal{ALCQI}_{reg}$ for recording other kinds of knowledge than those showed in the previous parts of this paper. Typical documentation systems, in practice, are expected to need much simpler kinds of knowledge, and to operate in a very specific scenario, allowing for a more reassuring complexity.

Indeed, in a consistent documentation system $\mathcal{S} = (D, O)$, there is an equivalence between D from the one hand, and O plus ν form the other, in the sense that D can be derived from O and ν, and conversely O and ν can be derived from D. This equivalence suggests an obvious way of managing a documentation system: curators add descriptions on the object-side, thereby expanding O, and while doing so also provide identifiers for the concepts that they create, thereby defining ν in an incremental way. Based on this input, the underlying software system can generate the corresponding doc-KB D, in a fully automated way, whenever D is needed. The user uses the doc-side only for inserting the knowledge that cannot be expressed in this way, namely:

- to assign descriptions to description identifiers, by making assertions on the role CAss;

- to create metadata records, by making assertions on the roles MRD and MRS.

Whenever this happens, the underlying software system generates the corresponding doc-assertions (as established by the translation function ϕ.)

This way of proceeding has the obvios advantage that D will always be fully consistent and "aligned" with O so as to satisfy the relation $O = \phi(D)$. Moreover, the only implicit documentation knowledge is that resulting from axiom 1; such knowledge can be made explicit by computing the composition of the involved role assertions, in a very simple and efficient way.

In sum, the complexity of maintaining and querying a documentation system is not affected by the basic documentation knowledge, expressing the connection between identifiers and identified descriptions. It is only affected by the kind of knowledge that is added on either side, for application purposes.

8. RELATED WORK

We will discuss related work in three main areas: the web, covering the digital library side, DLs and ontologies. Although there are approaches in all three areas that share some commonalities with our work, the notion of documentation system does not have any counterpart in the existing literature, to the best of our knowledge. As already pointed out, the present work also extends a previous attempt by the author at modelling documentation knowledge [10], whose expressive power is however limited to that of RDFS.

In the context of the web, Linked Data promote the usage of metadata records ("useful information" about a resource [2]) as web resources of their own right, identified via HTTP URIs. This basic mechanism is rapidly gaining ground, as Linked Data become increasingly popular[2]. There are nowadays millions of HTTP URIs that identify metadata records on the web, yet the potentialities of this idea are exploited only to a minimal extent, due to the lack of formal foundations. A somewhat more general approach is taken by OAI-ORE [8], which proposes entire RDF graphs, known as *Resources Maps,* as web resources, also identified via HTTP URIs. Resource Maps are more general than Linked Data graphs, because they can contain RDF triples about anything. The OAI-ORE vocabulary presents some overlaps with our doc-logic; in particular, the OAI-ORE property *Describes,* connecting a Resource Map to the aggregation described by it, can be seen as a special case of the CAss role. But this is the only relation between doc-logic and OAI-ORE, which is more focussed on the modelling of aggregations on the web.

In the DL literature, the work reported here relates to *reification,* generally understood as the method by which statements in a certain logic are represented as individuals in the same or in a different logic. In DL, reification means to translate a relation, into a concept whose instances are one-to-one with the relation tuples. Reification is used (all references are from [1]):

- to translate semantic data models into DL (sec. 4.3);

---

[2]For an overview of the status of Linked Data datasets, see http://linkeddata.org/

- to reason with expressive variants of number restrictions (sec. 5.4);

- to develop DLs with relations of arbitrary arity (sec. 5.7);

- to represent properties of properties (sec. 10.6.1).

All this work is not relevant to the present work, because we reify concepts and not relations.

Top ontologies offer a wide range of concepts and properties for modelling vocabularies and typically include concepts for representing both domain and documentation knowledge. For instance in CIDOC-CRM [5], a description would be modelled as an information object (class *E73 Information Object*), and its properties, such as authorship. Ontologies have a descriptive nature and aim at coverage, in that they aim at capturing *all* aspects of reality in a general, application independent way. As such, ontologies are not much concerned about computational aspects and therefore are not analyzed from the perspective of the present work.

9. CONCLUSIONS

We have presented documentation systems as constituents of digital libraries, using description logics to represent and reason about domain and documentation knowledge at the same time. In particular, the KB of a documentation system consists of an object-KB and of a doc-KB. In the object-TBox, it is possible to express domain ontologies as terminological axioms, whereas in the object-ABox it is possible to express metadata records as DL concept assertions. In the doc-ABox, it is possible to make assertions involving descriptions and the metadata records of the object-KB Abox, thereby representing documentation knowledge. As such, a documentation system significantly extends the expressive capabilities of current digital libraries, mostly based on RDF.

We have exemplified documentation systems by using a well-known DL, $\mathcal{ALCO}$, as object-logic, and by defining the corresponding doc-logic *alco.* Our interest is not in the doc-logic per sè, but in the simultaneous usage of the doc-logic and of its corresponding object-logic to the end of satisfying two distinct categories of users of a DL: end users, from the one hand, interested in the domain knowledge; and digital librarians or digital curators, from the other hand, interested in documenting their activity by expressing documentation knowledge.

Our treatment has covered the theoretical aspects of documentation systems, in the interest of clarity and well-foundedness. In particular, we have focussed on the syntax and the semantics of the doc-logic, and highlighted its computational properties. We have also shown how to re-use the existing DL proof theory in order to tackle consistency checking and query answering in a documentation system.

However, our treatment has not been *exclusively* theoretical. First, our considerations on how to tackle consistency checking and query answering make the existing DL technology (Protege, RACER, Pellet, etc.) applicable to the practical implementation of a documentation system. Moreover, the discussion in Sectionœ7 shows how a documentation system can be put at work in a practical digital library setting. Indeed, anyone familiar with OWL can use a documentation system without having to learn anything new.

10. REFERENCES

[1] Franz Baader, Diego Calvanese, Deborah L. McGuinness, Daniele Nardi, and Peter F. Patel-Schneider, editors. *The Description Logic Handbook: Theory, Implementation, and Applications.* Cambridge University Press, 2nd edition, 2003.

[2] Tim Berners-Lee. Linked data. `http://www.w3.org/DesignIssues/LinkedData.html`.

[3] Tim Berners-Lee, James Hendler, and Ora Lassila. The semantic web. *Scientific American Magazine*, May 2001.

[4] Dan Brickley and R.V. Guha. RDF vocabulary description language 1.0: RDF schema. W3C Recommendation, WWW Consortium, February 2004. `http://www.w3.org/TR/rdf-schema/`.

[5] Martin Doerr. The CIDOC conceptual reference model: An ontological approach to semantic interoperability of metadata. *AI Magazine*, 24(3):75–92, 2003.

[6] N. Guarino. Formal ontology in information systems. In *Proceedings of FOISÕ98*, pages 3–15. IOS Press, Amsterdam, 1998. Amended version.

[7] Graeme Hirst. Ontology and the lexicon. In Steffen Staab and Rudi Studer, editors, *Handbook on Ontologies*, pages 269–292. Springer, 2nd edition, 2009.

[8] Carl Lagoze and Herbert Van de Sompel (eds.). ORE user guide - Primer. Technical report, Open Archives Initiative. Object Reuse and Exchange, July 2008. `http://www.openarchives.org/ore/0.9/primer.html`.

[9] Frank Manola and Eric Miller. RDF Primer. W3C Recommendation, WWW Consortium, February 2004. `http://www.w3.org/TR/rdf-primer/`.

[10] Carlo Meghini, Nicolas Spyratos, Tsuyoshi Sugibuchi, and Jitao Yang. A model for digital libraries and its translation to rdf. *Journal on Data Semantics*, 2013.

[11] Boris Motik, Peter F. Patel-Schneider, and Bijan Parsia. OWL 2 Web Ontology Language structural specification and functional-style syntax (second edition). W3C recommendation, W3C, December 2012. `http://www.w3.org/TR/2012/REC-owl2-syntax-20121211/`.

[12] W3C OWL Working Group. OWL 2 Web Ontology Language document overview (second edition). W3C recommendation, W3C, December 2012. `http://www.w3.org/TR/owl2-overview/`.