UNIVERSITY
OF TRENTO

DEPARTMENT OF INFORMATION AND COMMUNICATION TECHNOLOGY

38050 Povo – Trento (Italy), Via Sommarive 14
http://www.dit.unitn.it

CONTEXTUALIZING ONTOLOGIES

Paolo Bouquet, Fausto Giunchiglia, Frank van Harmelen,
Luciano Serafini and Heiner Stuckenschmidt

February 2004

Technical Report # DIT-04-013

Also: to appear in the Journal of Web Semantics, 2004

Contextualizing Ontologies*

Paolo Bouquet[1,2], Fausto Giunchiglia[1,2], Frank van Harmelen[3], Luciano Serafini[1], and Heiner Stuckenschmidt[3]

[1] ITC-IRST, Trento, Italy
[2] DIT - University of Trento, Trento, Italy
[3] AI Department, Vrije Universiteit Amsterdam, The Netherlands

February 9, 2004

Abstract. Ontologies are *shared* models of a domain that encode a view which is common to a set of different parties. Contexts are *local* models that encode a party's subjective view of a domain. In this paper we show how ontologies can be contextualized, thus acquiring certain useful properties that a pure shared approach cannot provide. We say that an ontology is contextualized or, also, that it is a *contextual ontology*, when its contents are kept local, and therefore not shared with other ontologies, and mapped with the contents of other ontologies via explicit (context) mappings. The result is *Context OWL (C-OWL)*, a language whose syntax and semantics have been obtained by extending the OWL syntax and semantics to allow for the representation of contextual ontologies.

1 Introduction

The aim of the Semantic Web is to make information on the World Wide Web more accessible using machine-readable meta-data. In this context, the need for explicit models of semantic information (terminologies and background knowledge) in order to support information exchange has been widely acknowledged by the research community. Several different ways of describing information semantics have been proposed and used in applications. However we can distinguish two broad approaches which follow somehow opposite directions:

Ontologies are *shared* models of some domain that encode a view which is common to a set of different parties [19];

Contexts are *local* (where *local* is intended here to imply *not shared*) models that encode a party's view of a domain [14, 13, 12].

Thus, ontologies are best used in applications where the core problem is the use and management of common representations. Many applications have been developed, for instance in bio-informatics [10], or for knowledge management purposes inside organizations [8]. Contexts, instead, are best used in those applications where the core problem is the use and management of local and autonomous representations with a need for a limited and controlled form of globalization (or, using the terminology used in the context literature, maintaining *locality* still guaranteeing semantic *compatibility* among representations [12]). Examples of uses of contexts are the classifications of documents [6], distributed knowledge management [3], the development and integration of catalogs [11, 4], peer-to-peer applications with a large degree of *autonomy* of the peer nodes but still with a strong need of *coordination* [22] (with

* This paper is an extension of the paper "C-OWL - Contextualizing Ontologies", presented at the Second International Semantic Web Conference 2003

autonomy and coordination being the behavioral counterpart of the semantic need of locality and compatibility).

Contexts and ontologies have both strengths and weaknesses. It can be argued that the strengths of ontologies are the weaknesses of contexts and vice versa. On the one hand, the use of ontologies enables the parties to communicate and exchange information. Shared ontologies define a common understanding of specific terms, and thus make it possible to communicate between systems on a semantic level. On the weak side, ontologies can be used only as long as consensus about their contents is reached. Furthermore, building and maintaining (!) them may become arbitrarily hard, in particular in a very dynamic, open and distributed domain like the Web. On the other hand, contexts encode not shared interpretation schemas of individuals or groups of individuals. Contexts are easier to define and to maintain. They can be constructed with no consensus with the other parties, or only with the limited consensus which makes it possible to achieve the desired level of communication and only with the "relevant" parties. On the weak side, since contexts are local to parties, communication can be achieved only by constructing explicit mappings among the elements of the contexts of the involved parties; and extending the communication to new topics and/or new parties requires the explicit definition of new mappings.

Depending on their attitude, from an epistemological point of view, some people would argue that ontologies are all we need, while others would argue the exact contrary, namely that contexts are all we need. Our attitude in this paper is quite pragmatical. We believe that ontologies and contexts have both some advantages and that, therefore, they should be integrated in the representational infrastructure of the Semantic Web. Thus, on the one hand, the intended meaning of terms provided by parties which are willing to share information can be more easily captured with an ontology (or a set of shared ontologies). On the other hand, multiple ontologies (or sets or shared ontologies) which contain information which *should not* be integrated (an obvious example being information which is mutually inconsistent) should be contextualized. We say that an ontology is contextualized, or that it is a *contextual ontology*, when its contents are kept local (and therefore not shared with other ontologies) and are put in relation with the contents of other ontologies via explicit mappings.

Our approach in this paper is as follows. We take the notion of ontology as the core representation mechanism for representing information semantics. To this end, we start from the standard Web ontology language OWL [17]. Notice that from OWL we inherit the possibility to have shared ontologies. We show, providing some motivating examples, that OWL cannot model certain situations (Section 4). Finally, we provide an extension of OWL, that we call *Context OWL* (*C-OWL*), which allows us to deal with all the examples of Section 4. C-OWL integrates in a uniform way the, somehow orthogonal, key architectural features of contexts and ontologies and the consequent semantic level differences.

The main technical contributions of this paper are the following:

1. We provide a (somewhat synthetic) description of OWL and its semantics, restating Patel-Schneider and Hayes' semantics [19], in a formal framework more adequate to be extended (adapted) with a contextualized interpretation. These are the contents of Section 3.
2. We modify the OWL semantics to make it able to deal with the motivating examples reported in Section **??**. These are the contents of Section 5.

3. We define the C-OWL syntax by taking the OWL syntax and by adding *bridge rules*, which allow to relate, at the syntactic and at the semantic level, concepts, roles and individuals in different ontologies. We call a set of bridge rules between two ontologies a *context mapping*. Thus a *contextual ontology* is an OWL ontology embedded in a space of other OWL ontologies and related to them via context mappings. We define the C-OWL semantics by taking the modified OWL semantics, as defined in Section 5. These are the contents of Section 6.
4. Finally in Section 7 we show how C-OWL can be used for the alignment (i.e., weak integration) of a set of independently developed medical ontologies.

The semantics of C-OWL is obtained by modifying the OWL semantics [19] using the ideas and notions originally developed in [5], which is based on the semantics of context (the, so called, Local Models Semantics [13]). The general notion of bridge rules were originally defined in [15] and further studied in [14, 13, 21, 6, 5]. The bridge rules proposed in this paper were first defined in [7]. Finally the constructs for representing bridge rules have been taken from the context markup language CtxML [6].

2 Ontologies vs. Contexts, or globalize vs. localize

At the architectural level, the crucial difference between the notions of context and ontology is in how mappings among multiple models are constructed:

– in OWL, mappings are not part of the language. The ability of combining models is restricted to the import of complete models and to the use of the imported elements by direct reference. Via the import mechanism, a set of local models is *globalized* in a unique shared model (which, however, keeps track of the original distinctions). It is assumed that references to external statements are only made for statements from imported models, however, this is strictly speaking not required. As a consequence, mappings rather implicitly exist in terms of mutual use of statements across models. Further, there are two ways of treating external statements: we can either treat the referred statement as a single fact with no further implications that the one it directly encodes, or we can use the complete model containing the fact as additional knowledge. This latter view is the one adopted in OWL, and it is the one we will consider in the following.
– in context-based approaches, local models are kept *localized*. A limited and completely controlled form of globalization is obtained by using explicit mappings. In this approach, mappings are regarded as projections of a local representation onto another, and are first class modelling elements with a unique identity. In other words, also mappings are viewed as part of a local representation. This view makes it possible to have multiple alternative mappings between the same pair of contexts, and to define mappings in one direction that differ from the mappings in the opposite direction.

This different bias towards localization/globalization, and the consequent very different treatment of mappings lead to important semantic differences. OWL is mainly inspired by the Tarskian style semantics of propositional description logics. A model theoretic semantics is provided by mapping the elements of existing models into an abstract domain, where concepts are represented by sets, relation by sets of tuples and instances by elements of that domain. When reasoning is performed across different models, then these models are assumed to share the interpretation domain. Thus, as a consequence, the mappings between two models become part of the overall model and define constraints on the elements of the original two models.

The situation is quite different when we move to contexts. In the Local Models Semantics, each context uses a local set of models and a local domain of interpretation. Relations between these local interpretation domains are established by *domain relations* which explicitly codify how elements in one domain map into elements of the other domain. Domain relations are indexed by source and target domain, making them irreversible and non-transitive; and bridge rules modify only the target context, leaving the source unaffected.

3 A global semantics for OWL

According to [19], an OWL ontology is a set of annotated *axioms* and *facts*, plus import references to other ontologies. OWL ontologies can be referenced by means of a URI. Ontologies can also have annotations that can be used to record authorship and other information associated with an ontology. Since annotation directives have no effect on the semantics of OWL ontologies in the abstract syntax, we ignore them. We concentrate on the OWL-DL fragment of OWL. This language is equivalent to the SHOIQ(D+) DL, i.e., SHIQ(D+) extended with an equivalent of the oneOf constructor. The proposed framework can be restricted or generalized to OWL-lite and OWL-full, respectively.

Let I be a set of indexes, standing for a set of URI's of onotlogies. For instance I contains `http://www.w3.org/2002/07/owl`. Let also $\mathbb{C}$, $\mathbb{R}$ and $\mathbb{O}$ be the sets of strings that can used to denote concepts, roles and individuals respectively. The disjoint union of $\mathbb{C}$, $\mathbb{R}$ and $\mathbb{O}$ is denoted with $\mathbb{L}$.

Definition 1 (OWL Ontology). *An* OWL Ontology *(or simply an* ontology*) is a pair $\langle i, O_i\rangle$, where $i \in I$ and $O_i = \langle T_i, A_i\rangle$ where T and A are a T-box and an A-box respectively in the SHOIQ(D+) description logic on $\mathbb{L} \cup (I \times \mathbb{L})$. $\langle i, O_i\rangle$ is an ontology with index i.*

Suppose that $C, D, E, F \in \mathbb{C}$ and $r, s \in \mathbb{R}$. The following are examples of concepts that can appear in O_i.

$$C, \quad i\!:\!C, \quad C \sqcap D, \quad j\!:\!E, \quad C \sqcap (j\!:\!E), \quad \exists r.C \sqcup D, \quad \exists (j\!:\!s).C \sqcup (j\!:\!F) \tag{1}$$

Every expression occurring in O_i without an index is intended to be in the language defined by O_i, L_i. The expressions appearing in O_i with indexes j are supposed to be defined in O_j; therefore they appear in O_j without index or with the index j. We introduce the notions of *local language* and *foreign language*.

Definition 2 (Local language). *A* local concept*, w.r.t. i, is an element of $\mathbb{C}$ that appears in O_i either without indexes or with index equal to i.* Local roles *and* local individuals *are defined analogously. The set of local concepts, local roles, and local individuals w.r.t. i are denoted by $\mathbb{C}_i$, $\mathbb{R}_i$, and $\mathbb{O}_i$. The* local language *to i, $\mathbb{L}_i$, is the disjoint union of them.*

Local objects of a language L_i are also called *i-objects*. For notational convenience, in the following we always use the colon notation. Thus, for instance, that local concepts $C \in \mathbb{C}_i$ of an ontology O_i are written as $i : C$. A *foreign concept*, or equivalently a *non local concept*, w.r.t. $i \in I$, is a concept that appears in O_i but is defined in some ontology O_j. Foreign concepts are referred with the notation $j : c$. An analogous definition can be given for roles and individuals.

Definition 3 (Foreign language). *For any $j \neq i$, a j-*foreign concept *w.r.t. i is an element of $\mathbb{C}$ that appears in O_i with index j. j-*foreign roles *and j-*foreign individuals *are defined analogously. The j-foreign language w.r.t. i is the disjoint union of them.*

Among the concepts described in (1), C and D are local concepts w.r.t. i and r is a local role (w.r.t. i), while E and F are j-foreign concepts and s is a j-foreign role. By means of foreign concepts, roles and individuals, two ontologies can refer to the *same* semantic object defined in a third ontology.

Definition 4 (OWL space). *An* OWL space *is a family of ontologies $\{\langle i, O_i\rangle\}_{i\in I}$ such that every O_i is an ontology, and for each $i \neq j$, the j-foreign language of O_i is contained in the local language of O_j.*

Moving to semantics, the idea is now to restate the semantics in [19] making explicit reference to the notions of local and foreign language. This distinction, crucial for the work developed in the next section, is not made in [19].

The semantics for OWL spaces defined in [19] is based on the intuition that, in OWL, as in RDF, a data type denotes the set of data values that is the value space for the data type. Concepts denote sets of individuals. Properties relate individuals to other information, and are divided into two disjoint groups, data-valued properties and individual-valued properties. Data-valued properties relate individuals to data values; individual-valued properties relate individuals to other individuals.

In the following we assume that any domain we introduce (denoted by Δ possibly with indexes) contains the union of the value spaces of the OWL data types and Unicode strings.

Definition 5 (OWL interpretation [19]). *An* OWL interpretation *for the OWL space $\{\langle i, O_i\rangle\}_{i\in I}$, is a pair $\mathcal{I} = \langle \Delta^{\mathcal{I}}, (.)^{\mathcal{I}}\rangle$, where $\Delta^{\mathcal{I}}$, contains a non-empty set of objects (the resources) and $(.)^{\mathcal{I}}$ is a function such that*

1. *$\mathcal{I}(i, C) \subseteq \Delta^{\mathcal{I}}$ for any $i \in I$ and $C \in \mathbb{C}_i$;*
2. *$\mathcal{I}(i, r) \subseteq \Delta^{\mathcal{I}} \times \Delta^{\mathcal{I}}$ for any $i \in I$ and $r \in \mathbb{R}_i$;*
3. *$\mathcal{I}(i, o) \in D^{\mathcal{I}}$ for any $i \in I$ and $o \in \mathbb{O}_i$;*

Notice that $(.)^{\mathcal{I}}$ can be extended to all the complex descriptions of SHIQ(D+) as usual. Statements contained in the A-box and the T-box (i.e., facts and axioms) of an ontology O_i of an OWL space $\{\langle i, O_i\rangle\}_{i\in I}$ can be verified/falsified by an interpretation according the axioms written in [19].

We call the above interpretation, a *global interpretation*, to emphasize the fact that language is interpreted against a global domain. We call the overall approach, the *global semantics* approach to OWL.

Definition 6 (OWL axiom and fact satisfiability [19]). *Given an* OWL interpretation $\mathcal{I}$ *for* $\{\langle i, O_i\rangle\}_{i\in I}$, $\mathcal{I}$ *satisfies a fact or an axiom* ϕ *of the* O_i *according to the rules defined in the table "Interpretation of Axioms and Facts" of [19]. An OWL interpretation* $\mathcal{I}$ *satisfies an OWL space* $\{\langle i, O_i\rangle\}_{i\in I}$, *if* $\mathcal{I}$ *satisfies each axiom and fact of* O_i, *for any* i.

Notice that we do not give any interpretation of the possibility for O_i to import another ontology O_j. However, from the logical point of view, importing O_j into O_i can be thought of as duplicating all the statements of O_j in O_i.

4 Motivating Examples

We provide some examples which cannot be represented with the current syntax and semantics of OWL. These examples show the need to enrich ontologies with the capability to cope with:

1. *the directionality of information flow*: we need to keep track of the source and the target ontology of a specific piece of information;
2. *Local domains*: we need to give up the hypothesis that all ontologies are interpreted in a single global domain;
3. *Context mappings*: we need to be able to state that two elements (concepts, roles, individuals) of two ontologies, though being (extensionally) different, are *contextually* related, for instance because they both refer to the same object in the world.

Example 1 (Directionality). Consider two ontologies O_1 and O_2 and suppose that O_2 is an extension of O_1, i.e., O_2 imports O_1 and adds it some new axiom. Directionality is fulfilled if the axioms added to O_2 should not affect what is stated in O_1. Consider the case where O_1 contains the axioms $A \sqsubseteq B$ and $C \sqsubseteq D$; furthermore, suppose that O_2 contains the axiom $B \sqsubseteq C$. We would like to derive $A \sqsubseteq D$ in O_2 but not in O_1.

Let us see how the global semantics behaves in this case. Let $\{\langle 1, O_1\rangle, \langle 2, O_2\rangle\}$ be the OWL space containing O_1 and O_2. Let A, B, C, and D be 1 local concepts. Suppose that O_1 contains the axioms $A \sqsubseteq B$ and $C \sqsubseteq D$. Suppose that O_2 imports O_1, this implies that O_2 contains $1\!:\!A \sqsubseteq 1\!:\!B$ and $1\!:\!C \sqsubseteq 1\!:\!D$. Finally suppose that O_2 contains the extra axiom $1 : B \sqsubseteq 1 : C$. We have that any interpretation of $\{\langle 1, O_1\rangle, \langle 2, O_2\rangle\}$, should be such that $(1\!:\!A)^{\mathcal{I}} \subseteq (1\!:\!B)^{\mathcal{I}} \subseteq (1\!:\!C)^{\mathcal{I}} \subseteq (1\!:\!D)^{\mathcal{I}}$; and therefore $(1\!:\!A)^{\mathcal{I}} \subseteq (1\!:\!D)^{\mathcal{I}}$. This means that $1\!:\!A \sqsubseteq 1\!:\!D$ is a logical consequence of the statements contained in the OWL space and, therefore, that directionality is not fulfilled.

Example 2 (A special form of directionality: the propagation of inconsistency). Consider the previous example and suppose that O_2 contains also the following two facts: $1 : A(a)$ and $1 : \neg D(a)$. O_2 is inconsistent, but we want to avoid the propagation of inconsistency to O_1. However, this is not possible as the fact that there is no interpretation that satisfies the axioms in O_2, automatically implies that there is no interpretation for the whole OWL space, either.

Example 3 (Local domains). Consider the ontology O_{WCM} of a worldwide organization on car manufacturing. Suppose that O_{WCM} contains the "standard" description of a car with its components. Clearly such a domain should be abstract and general enough so that it could be used (imported) by a large set of users dealing with cars. O_{WCM} contains the concept car which is supposet to capture any possible car, not only the actual physical cars in circulation. O_{WCM} contains also a general axiom stating that a car has exactly one engine.

$$\mathsf{Car} \sqsubseteq (\geq 1)\mathsf{hasEngine} \sqcap (\leq 1)\mathsf{HasEngine} \tag{2}$$

Suppose that two car manufacturing companies, say Ferrari and Porche, decide to adopt the WCM standard and import it in their ontologies, O_{Ferrari} and O_{Porche}. The two companies customize the general ontology provided by WCM by adding the fact that the engine of a car is one of the engines they produce. Therefore, the following two axioms are added to the ontologies O_{Ferrari} and O_{Porche} respectively.

$$\mathsf{WCM}:\mathsf{car} \sqsubseteq \forall \mathsf{hasEngine}.\{\mathsf{F23}, \mathsf{F34i}\} \tag{3}$$

$$\mathsf{WCM}:\mathsf{car} \sqsubseteq \forall \mathsf{hasEngine}.\{\mathsf{P09}, \mathsf{P98i}\} \tag{4}$$

(3) states that, in the ontology O_{Ferrari}, a car has an F23 or an F45i engine (two Ferrari's engines). Similar interpretation is given to (4). Notice that the axioms above are supposed to have a local scope, i.e., they are supposed to be true only withing the ontology they are stated. However, from the semantical point of view, assuming global semantics implies that the effect of an axiom global. Indeed, according to the global semantics, any interpretation of the OWL space containing O_{WCM}, O_{Ferrari} and O_{Porche} is such that, either $(\mathsf{F23})^{\mathcal{I}_{\mathsf{Ferrari}}} = (\mathsf{P09l})^{\mathcal{I}_{\mathsf{Porche}}}$ or $(\mathsf{F34i})^{\mathcal{I}_{\mathsf{Ferrari}}} = (\mathsf{P09})^{\mathcal{I}_{\mathsf{Porche}}}$, which is not what we want as Ferrari does not produce Porche's engines and neither vice-versa. The main problem here is the diversity of the domains between O_{Ferrari} and O_{Porche}, and the fact that each of the two companies wants to reason in its own local domain, ignoring the fact that there are cars which engines different from the ones they produce.

Example 4 (Context mappings). Suppose we have an ontology O_{FIAT} describing cars from a manufacturing point of view, and a completely independent ontology O_{Sale} describing cars from a car vendor point of view. The two concepts of car defined in the two ontologies, (that can be referred by Sale : Car and FIAT : Car) are very different and it makes no sense for either ontology to import the concept of car from the other. The two concepts are not extensionally equivalent and the instances of FIAT:Car do not belong to Sale : Car and vice-versa. On the other hand the two concepts describe the same real-world class of objects from two different points of view, and there can be many reasons for wanting to integrate this information. For instance one might need to build a new concept which contains (some of) the information in Sale : Car and in FIAT : Car. This connection cannot be stated via OWL axioms, as, for instance

$$\mathsf{Sale}:\mathsf{Car} \equiv \mathsf{FIAT}:\mathsf{Car}$$

implies that

$$\mathsf{Car}^{\mathcal{I}_{\mathsf{Sale}}} = \mathsf{Car}^{\mathcal{I}_{\mathsf{FIAT}}}$$

i.e., that the two classes coincide at the instance level.

In this example, the problem is not only at the semantic level. As the following section will show, handling this example requires an extension of the OWL syntax.

5 A semantics for contextual ontologies

In this section we incrementally extend/modify the OWL global semantics, and in the last subsection, also its syntax, in order to be able to model the above examples.

5.1 Directionality

We modify the definition of interpretation given above according to the intuition described in [5]. The main idea is that we split a global interpretation into a family of (local) interpretations, one for each ontology. Furthermore, we allow for an ontology to be locally inconsistent, i.e., not to have a local interpretation. In this case we associate to O_i a special "interpretation" $\mathcal{H}$, called a *hole*, that verifies any set of axioms, possibly contradictory.

Definition 7 (Hole). *A* Hole *is a pair* $\langle \Delta^{\mathcal{H}}, (.)^{\mathcal{H}} \rangle$, *such that* $\Delta^{\mathcal{H}}$ *is a nonempty set and* $(.)^{\mathcal{H}}$ *is a function that maps every constant of* $\mathbb{O}_i$ *into an element of* $\Delta^{\mathcal{H}}$, *every concept of* $\mathbb{C}_i$ *in the whole* $\Delta^{\mathcal{H}}$ *and every role of* $\mathbb{R}_i$ *into the set* $\Delta^{\mathcal{H}} \times \Delta^{\mathcal{H}}$. $\mathcal{H}$ *is called a hole on* Δ^{H}.

Analogously to what done in [5], the function $(.)^{\mathcal{H}}$ can be extended to complex descriptions and complex roles in the obvious way.

Definition 8 (Satisfiability in a hole). $\mathcal{H}$ *satisfies all the axioms and facts, i.e., if* ϕ *is an axiom or a fact,* $\mathcal{H} \models \phi$.

Definition 9 (OWL interpretation with holes). *An* OWL interpretation with holes *for the OWL space* $\{\langle i, O_i \rangle\}_{i\in I}$, *is a family* $\mathcal{I} = \{\mathcal{I}_i\}_{i\in I}$, *where each* $\mathcal{I}_i = \langle \Delta^{\mathcal{I}_i}, (.)^{\mathcal{I}_i} \rangle$, *called the* local interpretation *of* O_i, *is either an interpretation of* L_i *on* $\Delta^{\mathcal{I}_i}$, *or it is a hole for* L_i *on* $\Delta^{\mathcal{I}_i}$, *and for all* $i \in I$, *each* $\Delta^{\mathcal{I}_i}$ *coincides and are equal to a set denoted by* $\Delta^{\mathcal{I}}$.

Each $(.)^{\mathcal{I}_i}$ can be extended in the usual way to interpret local descriptions. Foreign descriptions are interpreted by the combination of the different $(.)^{\mathcal{I}_i}$ for each $i \in I$. In particular for any concept, role or individual of the alphabet L_j, $(.)^{\mathcal{I}_i}$ can be extended to be the same as $(.)^{\mathcal{I}_j}$. Namely:

$$(j\!:\!x)^{\mathcal{I}_i} = (x)^{\mathcal{I}_j} \tag{5}$$

which can intuitively be read as, "the meaning of the j-foreign concept $j\!:\!x$ occurring in O_i is the same as the meaning of x occurring in O_j". Since all interpretations share the same domain, this semantics is well founded. Namely, the interpretation of j-foreign concepts in i are contained in the domain of i, $\Delta^{\mathcal{I}_i}$. In the following we give some examples of $(.)^{\mathcal{I}_i}$, for which we suppose that $C, D \in \mathbb{C}_i$ and $r \in \mathbb{R}_i$ and $D, F \in \mathbb{C}_j$ and $s \in \mathbb{R}_j$.

$$C^{\mathcal{I}_i} = \begin{cases} \text{Any subset of } \Delta^{\mathcal{I}_i} \text{ if } \mathcal{I}_i \neq \mathcal{H}_i \\ \Delta^{\mathcal{I}} \text{ otherwise} \end{cases}$$
$$(C \sqcap D)^{\mathcal{I}_i} = (C)^{\mathcal{I}_i} \cap (D)^{\mathcal{I}_i}$$
$$(C \sqcap j : E)^{\mathcal{I}_i} = (C)^{\mathcal{I}_i} \cap (E)^{\mathcal{I}_j}$$
$$(\neg C)^{\mathcal{I}_i} = \begin{cases} \Delta^{\mathcal{I}} \setminus (C)^{\mathcal{I}_i} \text{ If } \mathcal{I}_i \neq \mathcal{H}_i \\ \Delta^{\mathcal{I}} \text{ otherwise} \end{cases}$$
$$(j : E)^{\mathcal{I}_i} = (E)^{\mathcal{I}_j}$$
$$(\neg j : E)^{\mathcal{I}_i} = \begin{cases} \Delta^{\mathcal{I}} \setminus (E)^{\mathcal{I}_j} \text{ If } \mathcal{I}_i \neq \mathcal{H}_i \\ \Delta^{\mathcal{I}} \text{ otherwise} \end{cases}$$

Definition 10 (Axiom satisfiability). *Given an* OWL interpretation with holes*, $\mathcal{I}$ for $\{\langle i, O_i \rangle\}_{i \in I}$, $\mathcal{I}$ satisfies a fact or an axiom ϕ of the O_i, in symbols $\mathcal{I} \models i : \phi$ if $\mathcal{I}_i \models \phi$. An OWL interpretation $\mathcal{I}$ satisfies an OWL space $\{\langle i, O_i \rangle\}_{i \in I}$, if $\mathcal{I}$ satisfies each axiom and fact of O_i for each i.*

Notice that any global OWL interpretation $\mathcal{I}$, as defined in Definition 5, is a special case of an OWL interpretation with holes (Definition 9). This happens if every $\mathcal{I}_i$ is not a hole. So Definition 9 can be seen as an extension of Definition 5.

Let us see how holes affect satisfiability and ultimately how they allow to better model the intuitions behind OWL. A first effect of holes is that the same axiom can be satisfied in an ontology and not satisfied in another. Consider for instance the OWL interpretation with holes $\{\mathcal{I}_1, \mathcal{I}_2, \mathcal{H}_3\}$, where $\mathcal{I}_1$ and $\mathcal{I}_2$ are not holes. Suppose that $(A)^{\mathcal{I}_1} \not\subseteq (B)^{\mathcal{I}_2}$. Then we have that $1{:}A \sqsubseteq 2{:}B$ is not satisfied if it occurs in O_2, while it is satisfied if it occurs in O_3.

Example 5 (Examples 1 and 2 formalized). Consider the OWL interpretation with holes, $\mathcal{I} = \{\mathcal{I}_1, \mathcal{I}_2\}$ defined as follows

1. $\Delta^{\mathcal{I}_1} = \{a, b, c, d\}$, $A^{\mathcal{I}_1} = \{a\}$, $B^{\mathcal{I}_1} = \{a, b\}$, $C^{\mathcal{I}_1} = \{c\}$, $D^{\mathcal{I}_1} = \{c, d\}$,
2. $\Delta^{\mathcal{I}_2} = \{a, b, c, d\}$, and $\mathcal{I}_2 = \mathcal{H}_2$, i.e. $\mathcal{I}_2$ is a hole.

$\mathcal{I}$ is an interpretation for the OWL space containing O_1 and O_2, since

1. $\mathcal{I}_1 \models A \sqsubseteq B$, $\mathcal{I}_1 \models C \sqsubseteq D$, and $\mathcal{I}_1 \not\models A \sqsubseteq D$, by construction of $\mathcal{I}_1$,
2. $\mathcal{I}_2 \models 1{:}A \sqsubseteq 1{:}B$, $\mathcal{I}_2 \models 1{:}B \sqsubseteq 1{:}C$, and $\mathcal{I}_2 \models 1{:}C \sqsubseteq 1{:}D$, because $\mathcal{I}_2$ is a hole.

Notice that $\mathcal{I}$ is an interpretation that satisfies O_2 (i.e., $1 : A \sqsubseteq 1 : B$ $1 : B \sqsubseteq 1 : C$, and $1{:}C \sqsubseteq 1{:}D$)], without making $A \sqsubseteq D$ true in O_1.

To formalize Example 2, we consider the same interpretation as above. This interpretation satisfies any axiom in O_2 ($\mathcal{I}_2$ is a hole) still keeping O_1 consistent ($\mathcal{I}_1$ is an interpretation which is not a hole and which satisfies O_1).

5.2 Local domains

The OWL semantics described in the previous section assumes the existence of a unique shared domain, namely, that each ontology describes the properties of the whole universe. In many cases this is not true as, for instance, an ontology on cars is not supposed to speak about medicines, or food. The idea here is to associate to each ontology a *local domain*. Local domains may overlap as we have to cope with the case where two ontologies refer to the same object.

Definition 11 (OWL interpretation with local domains). *An* OWL interpretation with local domains *for the OWL space* $\{\langle i, O_i\rangle\}_{i\in I}$, *is a family* $\mathcal{I} = \{\mathcal{I}_i\}_{i\in I}$, *where each* $\mathcal{I}_i = \left\langle \Delta^{\mathcal{I}_i}, (.)^{\mathcal{I}_i}\right\rangle$, *called the* local interpretation *of* O_i, *is either an interpretation of* L_i *on* $\Delta^{\mathcal{I}_i}$, *or a hole.*

Definition 11 is obtained from Definition 9 simply by dropping the restriction on domain equality. The interpretation $(.)^{\mathcal{I}_i}$ is extended to complex concepts, roles, and individuals, in the usual way. We have to take care, however, that j-foreign concepts, roles, and individuals used in O_i could be interpreted (by the local interpretation $\mathcal{I}_j$) in a (set of) object(s) which are not in the local domain $\Delta^{\mathcal{I}_i}$. Indeed, to deal with this problem, we have to impose that any expression occurring in O_i should be interpretable in the local domain $\Delta^{\mathcal{I}_i}$. As a consequence, we restrict the interpretation of any foreign concept $C \in \mathbb{C}_j$, any foreign role $r \in \mathbb{R}_j$ and any foreign individual $a \in \mathbb{O}_j$ as follows:

1. $(j\!:\!C)^{\mathcal{I}_i} = (C)^{\mathcal{I}_j} \cap \Delta^{\mathcal{I}_i}$
2. $(j\!:\!r)^{\mathcal{I}_i} = (r)^{\mathcal{I}_j} \cap (\Delta^{\mathcal{I}} \times \Delta^{\mathcal{I}})$
3. $(j\!:\!a)^{\mathcal{I}_i} = (a)^{\mathcal{I}_j}$

Notice that point 3 above implicitly imposes that if a j-foreign constant $j\!:\!a$ is used in the ontology O_i, then its interpretation in j, i.e., $a^{\mathcal{I}_j}$, must be contained in the domain $\Delta^{\mathcal{I}_i}$. Let us now see how we can deal with Example 3.

Example 6 (Example 3 formalized). Consider the OWL interpretation with local domains, $\mathcal{I} = \{\mathcal{I}_{\mathsf{WCM}}, \mathcal{I}_{\mathsf{Ferrari}}, \mathcal{I}_{\mathsf{Porche}}\}$ for the OWL space containing O_{WCM}, O_{Ferrari}, and O_{Porche}. Suppose that Δ_{WCM} contains four individuals $c_1, \ldots, c_4$ for cars and four individuals $e_1, \ldots, e_4$ for engines, with $\mathsf{hasEngine}^{\mathcal{I}_{\mathsf{WCM}}} = \{\langle c_1, e_1\rangle, \ldots, \langle c_4, e_4\rangle\}$. Let $\Delta_{\mathsf{Ferrari}} = \{c_1, c_2, e_1, e_2\}$ and $\Delta_{\mathsf{Porche}} = \{c_3, c_4, e_3, e_4\}$. be the local domains for O_{Ferrari} and O_{Porche} respectively. Suppose that $\mathcal{I}_{\mathsf{Ferrari}}$ interprets $\mathsf{F23}$ and $\mathsf{F34i}$ in e_1 and e_2 respectively, and that $\mathcal{I}_{\mathsf{Porche}}$ interprets $\mathsf{P09}$ and $\mathsf{P98i}$ in e_3 and e_4 respectively.

This OWL interpretation with local domains satisfies all the axioms (2), as in $\mathcal{I}_{\mathsf{WCM}}$ a car has only one engine; it satisfies axioms (3) since the interpretation of $\mathsf{car} : \mathsf{WCM}$ in O_{Ferrari} is restricted to be $\{c_1, c_2\}$ whose engine is a ferrari engines. Analogously this OWL interpretation satisfies (4). Notice however that Ferrari's engines are disjoint from Porche's engines.

5.3 Context mappings

We have concepts, roles and individuals local to different ontologies and domains of interpretation. A context mapping allows us to state that a certain property holds between elements of two different ontologies. Thus, for instance, in Example 4, one possible mapping could allow us to say that the class Car in the ontology O_{FIAT} contains the same cars as (or, as we say, is contextually equivalent to) the class of Car defined in the ontology O_{Sale}. As from Example 4 this cannot be done via local axioms within an ontology.

The basic notion towards the definition of context mappings are *bridge rules.*

Definition 12 (Bridge rules). *A* bridge rule from i to j is a statement of one of the four following forms,

$$i:x \xrightarrow{\sqsubseteq} j:y, \quad i:x \xrightarrow{\sqsupseteq} j:y, \quad i:x \xrightarrow{\equiv} j:y, \quad i:x \xrightarrow{\perp} j:y, \quad i:x \xrightarrow{*} j:y,$$

where x and y are either concepts, or individuals, or roles of the languages L_i and L_j respectively.

A mapping between two ontologies is a set of bridge rules between them.

Definition 13 (Mapping). *Given a OWL space $\{\langle i, O_i\rangle\}_{i\in I}$ a* mapping *M_{ij} from O_i to O_j is a set of bridge rules from O_i to O_j, for some $i, j \in I$.*

Mappings are directional, i.e., M_{ij} is not the inverse of M_{ji}. A mapping M_{ij} might be empty. This represents the impossibility for O_j to interpret any i-foreign concept into some local concept. Dually M_{ij} might be a set of bridge rules of the form $i:x \xrightarrow{\equiv} j:y$ for any element x (concept, role, and individual) of O_i. This represents the operation of mapping all of O_i into an equivalent subset of O_j. If this subset is O_j itself then this becomes the contextual mapping version of the OWL import operation. However, notice that importing O_i into O_j is not the same as mapping O_i to O_j with M_{ij}. In both cases information goes from i to j. The difference is that, in the former case, O_j duplicates the information of i- foreign elements without any change, while, in the latter, O_j translates (via the mapping M_{ij}) the semantics of O_i into its internal (local) semantics.

Definition 14 (Context space). *A context space is a pair composed of an OWL space $\{\langle i, O_i\rangle\}_{i\in I}$ and a family $\{M_{ij}\}_{i,j\in I}$ of mappings from i to j, for each pair $i, j \in I$.*

To give the semantics of context mappings we extend the definition of OWL interpretation with local domains with the notion of *domain relation.* A domain relation $r_{ij} \subseteq \Delta^{\mathcal{I}_i} \times \Delta^{\mathcal{I}_j}$ states, for each element in $\Delta^{\mathcal{I}_i}$ to which element in $\Delta^{\mathcal{I}_j}$ it corresponds to. The semantics for bridge rules from i to j can then be given with respect to r_{ij}.

Definition 15 (Interpretation for context spaces). *An interpretation for a context space $\langle\{\langle i, O_i\rangle\}_{i\in I}, \{M_{ij}\}_{ij\in I}\rangle$ is composed of a pair $\langle\mathcal{I}, \{r_{ij}\}_{i,j\in I}\rangle$. where $\mathcal{I}$ is an OWL interpretation with holes and local domains of $\{\langle i, O_i\rangle\}_{i\in I}$ and r_{ij}, the domain relation from i to j, is a subset of $\Delta^{\mathcal{I}_i} \times \Delta^{\mathcal{I}_j}$.*

Definition 16 (Satisfiability of bridge rules[1]).

1. $\mathfrak{I} \models i:x \xrightarrow{\sqsubseteq} j:y$ *if* $r_{ij}(x^{\mathcal{I}_i}) \subseteq y^{\mathcal{I}_j}$;
2. $\mathfrak{I} \models i:x \xrightarrow{\sqsupseteq} j:y$ *if* $r_{ij}(x^{\mathcal{I}_i}) \supseteq y^{\mathcal{I}_j}$;
3. $\mathfrak{I} \models i:x \xrightarrow{\equiv} j:y$ *if* $r_{ij}(x^{\mathcal{I}_i}) = y^{\mathcal{I}_j}$;
4. $\mathfrak{I} \models i:x \xrightarrow{\perp} j:y$ *if* $r_{ij}(x^{\mathcal{I}_i}) \cap y^{\mathcal{I}_j} = \emptyset$;
5. $\mathfrak{I} \models i:x \xrightarrow{*} j:y$ $r_{ij}(x^{\mathcal{I}_i}) \cap y^{\mathcal{I}_j} \neq \emptyset$;

A interpretation for a context space is a model for it if all the bridge rules are satisfied.

When x and y are concepts, say C and D, the intuitive reading of $i:C \xrightarrow{\sqsubseteq} j:D$, is that the i-local concept C is more specific than the j-concept D. An analogous reading can be given to $i:C \xrightarrow{\sqsupseteq} j:D$. The intuitive reading of $i:C \xrightarrow{\perp} j:D$ is that C is disjoint from D. Finally, the intuitive reading of $i:C \xrightarrow{*} j:D$ is that C and D are two concepts which are compatible. When x and y are individuals, then $i:x \xrightarrow{\sqsubseteq} j:y$ states that y is a more abstract representation of the object represented by x in i (intuitively, there might be more than one x's corresponding to the same y) Vive-versa $i:x \xrightarrow{\sqsupseteq} j:y$ states that y is a less abstract (more concrete) representation of the object represented by x in i (intuitively there might be more than one y's corresponding to the same x). $i:x \xrightarrow{\equiv} j:y$ states that x and y are at the same level of abstraction. Notice that, we add $i:a \xrightarrow{\equiv} j:a$ for any individual a of Δ_i and Δ_j we reduce to the case of OWL interpretation with holes and local domains). $i:x \xrightarrow{\perp} j:y$ states that x and y denotes completely unrelated objects. While $i:x \xrightarrow{*} j:y$ states that x and y might be related.

Example 7 (Example 4 and 3 formalized). The fact that SaleCar describes the *same* set of objects from two different points of view, can be captured by asserting the bridge rule:

$$\mathsf{Sale}:\mathsf{Car} \xrightarrow{\equiv} \mathsf{FIAT}:\mathsf{Car} \tag{6}$$

The domain relation from O_{Sale} to O_{FIAT} of any contextual interpretation satisfying (6) will be such that $r_{ij}(\mathsf{Car})^{\mathcal{I}_{\mathsf{Sale}}} = (\mathsf{Car})^{\mathcal{I}_{\mathsf{FIAT}}}$.

6 C-OWL: Extending OWL

In the previous sections we showed how certain requirements with respect to a contextual representation, in particular local domains and directionality can be achieved by a modification of the OWL semantics keeping its syntax unchanged. This allows us to define Context OWL as a strict extension of the OWL standard This minimal invasive approach guarantees a wide applicability of the model proposed here. In fact we can create an OWL space by defining mappings between already existing ontologies on the web. What is left to be done is to define an appropriate language for representing mappings between OWL ontologies along the ideas presented in the previous section. C-OWL can therefore be straightforwardly obtained from

[1] In this definition, to be more homogeneous, we consider the interpretations of individuals to be sets containing a single object rather than the object itself.

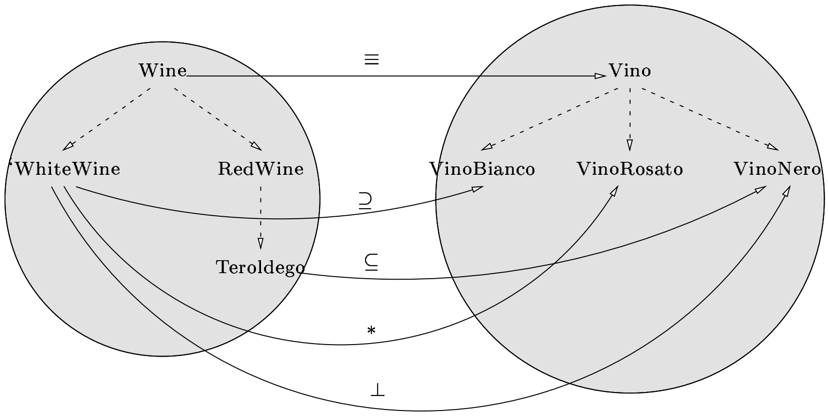

Fig. 1. A C-OWL mapping from the ontology "wine" to the ontology "vino".

CtxML by substituting the language for representing contexts in item 1 with OWL, and by keeping item 2 unchanged. As a consequence, C-OWL has the full representational power of OWL when we boil down to using ontologies, and the full representational power of CtxML when we boil down to using contextual information. The further nice property of C-OWL is that the two components are completely orthogonal and one can use the ontology or the contextual component in a totally independent manner.

In this section we define an RDF-based syntax for such mappings. We introduce the semantics using an example, explain the different parts of the specification and define an RDF schema for the mapping representation.

The philosophy of C-OWL is to treat mappings as first class and to represent them independently from the ontologies they connect. There are a couple of advantages of this approach. From a syntactic point of view, the advantage is that we can define a language for specifying mappings independently from the OWL syntax specification. the resulting language will refer to elements of the OWL specification without extending it.

Figure 1 shows an example mapping of two ontologies about wines. In order to represent this mapping we have to capture the following aspects:

- a unique identifier for referring to the mapping
- a reference to the source ontology
- a reference to the target ontology
- a set of bridge rules relating classes from the two ontologies, each described by
  - (a reference to) the source concept
  - (a reference to) the target concept
  - the type of the bridge rule, which is one of $\equiv$ $\sqsubseteq$, $\sqsupseteq$, $\perp$, $*$

Figure 2 shows an RDF-based representation of these elements. We use a resource of the type `cowl:Mapping` as a root element of the description. This resource is linked to two OWL models using the properties `sourceOntology` and `targetOntology`. The ontologies are represented by reference to their namespace. Further, the resource representing the overall

```xml
<?xml version="1.0" encoding="UTF-8"?>
<rdf:RDF
    xmlns:rdf="http://www.w3.org/1999/02/22-rdf-syntax-ns#"
    xmlns:rdfs="http://www.w3.org/2000/01/rdf-schema#"
    xmlns:cowl="http://www.cowl.org/"
    xmlns:owl="http://www.w3.org/2002/07/owl#">
    <cowl:Mapping rdf:ID="myMapping">
        <rdfs:comment>Example Mapping for Web Semantics Journal Paper</rdfs:comment>
        <cowl:sourceOntology>
            <owl:Ontology rdf:about="http://www.example.org/wine.owl"/>
        </cowl:sourceOntology>
        <cowl:targetOntology>
            <owl:Ontology rdf:about="http://www.example.org/vino.owl"/>
        </cowl:targetOntology>
        <cowl:bridgeRule>
            <cowl:Equivalent>
                <cowl:source>
                    <owl:Class rdf:about="http://www.example.org/wine.owl#wine"/>
                </cowl:source>
                <cowl:target>
                    <owl:Class rdf:about="http://www.example.org/vino.owl#vino"/>
                </cowl:target>
            </cowl:Equivalent>
        </cowl:bridgeRule>
        <cowl:bridgeRule>
            <cowl:Onto>
                <cowl:source>
                    <owl:Class rdf:about="http://www.example.org/wine.owl#RedWine"/>
                </cowl:source>
                <cowl:target>
                    <owl:Class rdf:about="http://www.example.org/vino.owl#VinoRosso"/>
                </cowl:target>
            </cowl:Onto>
        </cowl:bridgeRule>
        <cowl:bridgeRule>
            <cowl:Into>
                <cowl:source>
                    <owl:Class rdf:about="http://www.example.org/wine.owl#Teroldego"/>
                </cowl:source>
                <cowl:target>
                    <owl:Class rdf:about="http://www.example.org/vino.owl#VinoRosso"/>
                </cowl:target>
            </cowl:Into>
        </cowl:bridgeRule>
        <cowl:bridgeRule>
            <cowl:Compatible>
                <cowl:source>
                    <owl:Class rdf:about="http://www.example.org/wine.owl#WhiteWine"/>
                </cowl:source>
                <cowl:target>
                    <owl:Class rdf:about="http://www.example.org/vino.owl#Passito"/>
                </cowl:target>
            </cowl:Compatible>
        </cowl:bridgeRule>
        <cowl:bridgeRule>
            <cowl:Incompatible>
                <cowl:source>
                    <owl:Class rdf:about="http://www.example.org/wine.owl#WhiteWine"/>
                </cowl:source>
                <cowl:target>
                    <owl:Class rdf:about="http://www.example.org/vino.owl#VinoNero"/>
                </cowl:target>
            </cowl:Incompatible>
        </cowl:bridgeRule>
    </cowl:Mapping>
</rdf:RDF>
```

Fig. 2. Specification of the Mappings from figure 1

mapping is linked to a number of resources through the `cowl:bridgeRule` property. These resources represent the individual rules in the mappings and can be of type `cowl:Equivalent`, `cowl:Into`, `cowl:Onto`, `cowl:Incompatible` or `cowl:Compatible` each representing one of the types mentioned above. Each of the resources representing a bridge rule is linked to an OWL class from the target ontology through the `cowl:source` and to a class from the target ontology by the `cowl:target` property. The classes can be represented by a reference to the corresponding resource in the ontology definition but it can also be a complex OWL class definition that uses elements from the respective ontology. In this way we can represent complex mappings that go beyond semantic relations between class names. We have defined an RDF schema for the mapping representation. This schema is shown in figure 3

7 Aligning Medical Ontologies with C-OWL

The need for terminology integration has been widely recognized in the medical area leading to a number of efforts for defining standardized terminologies. It is, however, also acknowledged by the literature, that the creation of a single universal terminology for the medical domain is neither possible nor beneficial, because different tasks and viewpoints require different, often incompatible conceptual choices [9]. As a result a number of communities of practice have been evolved that commit to one of the proposed standards. This situation demands for a weak for of integration, also referred to as alignment in order to be able to exchange information between the different communities.

The notion of contextualized ontologies can provide such an alignment by allowing the coexistence of different, even in mutually inconsistent models that are connected by semantic mappings. As discussed above, the nature of the proposed semantic mappings satisfies the requirements of the medical domain, because they do not require any changes to the connected ontologies and do not create logical inconsistency even if the models are incompatible.

7.1 (Bio-)Medical Ontologies

In the medical area a lot of work has been done on the definition and standardization of terminologies[2]. The result of these efforts is a large number of medical terminologies and classifications. The complexity of the terminologies used in medicine and the strong need for quality control has also lead to the development of ontologies that feature complex concept definition (compare [16] for a discussion of the required expressiveness). Some of these ontologies are available in OWL and can be seen as the first OWL applications that have a use in real life applications. We briefly introduce three medical ontologies that are available in OWL.

Galen The Motivation for the GALEN project [20] is the difficulty in exchanging clinical data between different persons and organizations due to the heterogeneity of the terminology used. As a result of the project, the GALEN Coding Reference model has been developed. This reference model is an ontology that covers general medical terms, relations between

[2] See e.g. `http://www.medinf.muluebeck.de/~ingenerf/terminology/Index.html` for a collection of standards

```
<?xml version="1.0" encoding="UTF-8"?>
<rdf:RDF xmlns:rdf="http://www.w3.org/1999/02/22-rdf-syntax-ns#"
         xmlns:rdfs="http://www.w3.org/2000/01/rdf-schema#"
         xmlns:owl ="http://www.w3.org/2002/07/owl#">
    <rdfs:Class rdf:about="Mapping"/>
    <rdfs:Class rdf:about="Correspondence"/>
    <rdfs:Class rdf:about="Equivalence">
        <rdfs:subClassOf rdf:resource="#Correspondence"/>
    </rdfs:Class>
    <rdfs:Class rdf:about="Onto">
        <rdfs:subClassOf rdf:resource="#Correspondence"/>
    </rdfs:Class>
    <rdfs:Class rdf:about="Into">
        <rdfs:subClassOf rdf:resource="#Correspondence"/>
    </rdfs:Class>
    <rdfs:Class rdf:about="Compatible">
        <rdfs:subClassOf rdf:resource="#Correspondence"/>
    </rdfs:Class>
    <rdfs:Class rdf:about="Incompatible">
        <rdfs:subClassOf rdf:resource="#Correspondence"/>
    </rdfs:Class>
    <rdf:Property rdf:about="sourceOntology">
        <rdfs:domain rdf:resource="#Mapping"/>
        <rdfs:range rdf:resource="owl:Ontology"/>
    </rdf:Property>
    <rdf:Property rdf:about="targetOntology">
        <rdfs:domain rdf:resource="#Mapping"/>
        <rdfs:range rdf:resource="owl:Ontology"/>
    </rdf:Property>
    <rdf:Property rdf:about="bridgeRule">
        <rdfs:domain rdf:resource="#Mapping"/>
        <rdfs:range rdf:resource="#Correspondence"/>
    </rdf:Property>
    <rdf:Property rdf:about="source">
        <rdfs:domain rdf:resource="#Correspondence"/>
        <rdfs:range rdf:resource="owl:Class"/>
    </rdf:Property>
    <rdf:Property rdf:about="target">
        <rdfs:domain rdf:resource="#Correspondence"/>
        <rdfs:range rdf:resource="owl:Class"/>
    </rdf:Property>
</rdf:RDF>
```

Fig. 3. RDF schema defining the Extensions to OWL

those terms as well as complex concepts that are defined using basic terms and relations. We used an OWL version of the GALEN model that contains about 3100 classes and about 400 relations.

Tambis The aim of the Tambis [1] (Transparent Access to Bioinformatics Information Sources) is to provide an infrastructure that allows researchers in Bioinformatics to access multiple sources of biomedical resources in a single interface. In order to achieve this functionality, the project has developed the Tambis Ontology, which is an explicit representation of biomedical terminology. The complete version of Tambis contains about 1800 terms. The DAML+OIL version we used in the case study actually contains a subset of the complete ontology. It contains about 450 concepts and 120 Relations.

UMLS The Unified Medical Language System UMLS [18] is an attempt to integrate different medical terminologies and to provide a unified terminology that can be used across multiple medical information sources. Examples of medical terminologies that ave been integrated in UMLS are MeSH and SNOWMED. In our case study, we used the UMLS semantic network. The corresponding model that is available as OWL file contains 134 semantic types organized in a hierarchy as well as 54 relations between them with associated domain and range restrictions.

7.2 Alignment Scenario

C-OWL and especially its formal semantics provides us with several possibilities concerning the alignment of the medical ontologies mentioned above. We assume that the goal is to establish a connection between the Tambis and the GALEN ontology in such a way that the two models with their different focus supplement each other. The first option is to directly link the two ontologies by defining appropriate bridge rules which formalizes the semantic relation between concepts in the two ontologies. These bridge rules can be represented using the syntax described in the previous section and stored in separated files that can be used by a third parties. A second option for aligning Tambis and GALEN is based on a third, already existing, more general model of the domain (UMLS in this case). In this setting the relation between Tambis and GALEN can be *logically inferred* from the relations between each single ontology and the more general ontology UMLS as shown in Figure 4

Being the result of an integration of different medical terminologies(compare [2]) the UMLS semantic network is such a general model, that we can assume it as a general medical ontology that covers most of the content of Tambis, GALEN and also other prospective ontologies that we might want to align. Its important to notice that the fact that UMLS completely covers GALEN and Tambis is not a strong requirement, as a partial coverage does not prevents us to define partial alignment.

In order to explore the use of C-OWL for the alignment of medical ontologies, we conducted a small case study in aligning the ontologies mentioned above using the UMLS semantic network as a central terminology. We investigated the upper parts of the ontologies and identified areas with a sufficient overlap. Such an overlap between all three models exists with respect to the following three areas:

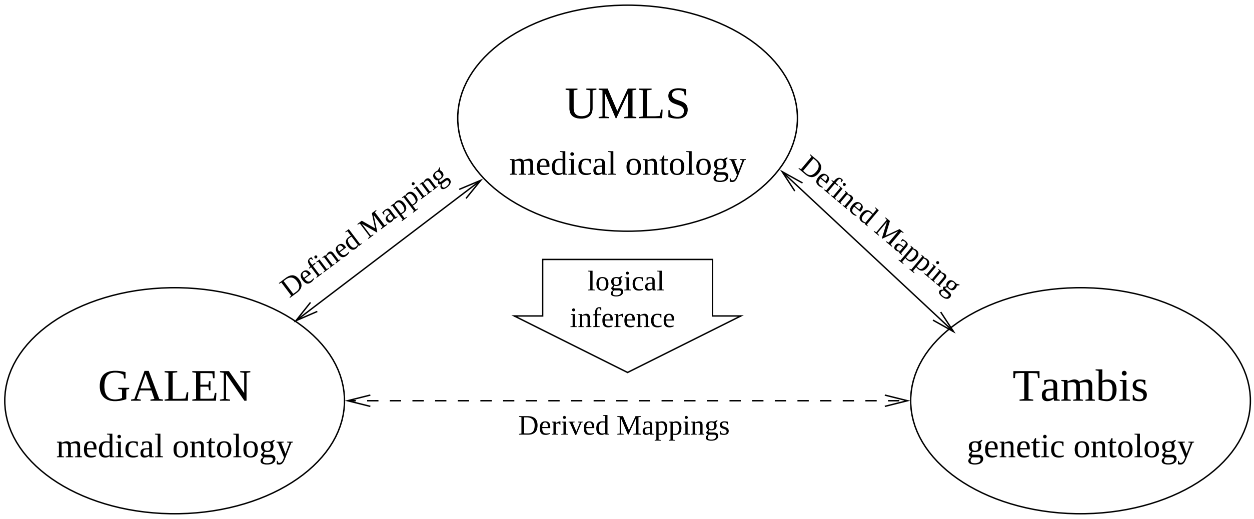

Fig. 4. Indirect Alignment of Tambis and GALEN using UMLS

Processes: Different physiological, biological and chemical processes related to the functioning of the human body and to the treatment of malfunctions.

Substances: Substances involved in physiological processes including chemical, biological and physical substances.

Structures: Objects and object assemblies that form the human body or parts of it. Further, structures used in the treatment of diseases.

We analyzed the three models with respect to these three topics. Based on the comparison of the three models, we define mappings between Tambis and GALEN and the UMLS terminology. These mappings consist of sets of bridge rules each connecting single concepts or concept expressions.

In the following, we discuss the ability of C-OWL to reason about the defined mappings using examples from the substances topic. We describe inferred knowledge about the mappings in terms of detected inconsistencies and derived semantic relations between the two ontologies.

7.3 Examples from the Alignment

GALEN contains the notion of a generalized substance which is a notion of substance that subsumes substances in a physical sense and energy making it more general than the notion of substance in UMLS

$$\mathsf{GeneralisedSubstance} \xleftrightarrow{\sqsupseteq} \mathsf{Substance}$$

The actual notion of substance as defined in GALEN is not as we might expect equivalent to the notion of substance in UMLS, because it also contains some notions that are found under anatomical structures in UMLS. We can, however, state that the GALEN notion of substance is more specific than the union of substances and anatomical structures in UMLS.

$$\mathsf{Substance} \xleftrightarrow{\sqsubseteq} \mathsf{Substance} \sqcup \mathsf{Anatomical\_Structure}$$

The next GALEN concept that also occurs in UMLS but has a slightly different meaning is the notion of body substance. The difference is illustrated in the fact that it also covers the notion of tissue which is found under anatomical structures in UMLS. We conclude that the notion of body substance in GALEN in a broader one than in UMLS.

$$\mathsf{BodySubstance} \xleftrightarrow{\sqsupseteq} \mathsf{Body\_Substance}$$

The other main class of substances mentioned in GALEN are chemical substances. Looking at the things contained under this notion, we conclude that it is equivalent to the notion of chemical in UMLS.

$$\mathsf{ChemicalSubstance} \xleftrightarrow{\equiv} \mathsf{Chemical}$$

We can also find the correspondences to the distinction between elementary and complex chemicals made by GALEN in UMLS. Elementary chemicals are a special case of the UMLS concept of elements ion or isotope.

$$\mathsf{ElemetaryChemical} \xleftrightarrow{\sqsubseteq} \mathsf{Element\_Ion\_or\_Isotop}$$

Complex chemicals contain all kinds of chemical substances sometimes viewed structurally, sometimes functionally. Therefore, we cannot related this concept to one of these views taken by UMLS. We also notice that there are notions of complex chemicals in GALEN that do not occur under chemicals in UMLS - e.g. Drugs that related to the concept of clinical drug classified under manufactured objects.

$$\mathsf{Drug} \xleftrightarrow{\equiv} \mathsf{Clinical\_Drug}$$

Further, the UMLS views on chemicals also contain elementary chemicals. Consequently, we can only define the notion of complex chemical to be compatible with the union of the two views in UMLS

$$\mathsf{ComplexChemical} \xleftrightarrow{*} \mathsf{Chemical\_Viewed\_Structurally} \sqcup \mathsf{Chemical\_Viewed\_Functional}$$

On the level of more concrete chemical notions we find a number of correspondences mentioned in the following. Named hormones are equivalent to hormones in UMLS

$$\mathsf{NAMEDHormone} \xleftrightarrow{\equiv} \mathsf{Hormone}$$

Proteins are more specific than amino acids, peptides or proteins.

$$\mathsf{Protein} \xleftrightarrow{\sqsubseteq} \mathsf{Amino\_Acid\_Peptide\_or\_Protein}$$

The notions of lipid and of carbohydrate are the same in the two models

$$\mathsf{Lipid} \xleftrightarrow{\equiv} \mathsf{Lipid}$$

$$\mathsf{Carbohydrate} \xleftrightarrow{\equiv} \mathsf{Carbohydrate}$$

There is an overlap between the notion of acid in GALEN and the concepts amino acid, peptide or protein and Nucleic acid , nucleosid or protein in UMLS.

$$\mathsf{Acid} \xleftrightarrow{*} \mathsf{Amino\_Acid\_Peptide\_or\_Protein} \sqcup \mathsf{Nucleic\_Acid\_Nucleosid\_or\_Protein}$$

Finally metals can be defined to be a special case of inorganic chemicals.

$$\mathsf{Metal} \xleftrightarrow{\sqsubseteq} \mathsf{Inorganic\_Chemical}$$

In summary, we were able to find a lot of correspondences on the level of groups of chemicals. While the models disagreed on the higher level structuring of substances, they shared a lot of more concrete concepts. As a consequence, we found a number of equivalence and subsumption relationships between substances at a lower level while at the more general level, we often had to use weak relations or link to very general concepts.

7.4 Benefits of using C-OWL

In the experiment, we defined mappings in a ad-hoc rather than a systematic fashion. Such an ad hoc approach for defining mappings bears the risk of inconsistency and in completeness. We cannot prevent the definition of inconsistent or incomplete mappings, but the semantics of C-OWL can be used to verify and extend a defined mapping in order to detect inconsistencies and implied mappings. In the following we give examples of the use of the C-OWL semantics to verify and extend the mappings between the substance information in the different medical ontologies.

Verification of Mappings A mapping can become inconsistent if two classes who are known to overlap, e.g. because they are subclasses of each other, link to disjoint concepts in another model. An example of this situation can be found in the substance related part of the alignment. Figure 5 shows the situation. On the right hand side the extensions of the UMLS concept chemical substances and some of its subclasses are sketched. UMLS distinguishes between chemical from a structural and a functional view. In the case where these two views are defined to be disjoint (one can either take a structural or a functional view but not both) we get an inconsistency with the mappings defined for the Tambis ontology, because the mappings claims that the image of the concept chemical is exactly the extension of the structural view. At the same time, we claim that the image of enzyme which is a subclass of chemical is exactly the extension of the UMLS concept Enzyme which is classified under the functional view on chemicals in UMLS and therefore disjoint from the structural view. This however is now possible in the C-OWL semantics as the image of enzyme is a subset of the image of chemical by definition.

This ability to detect inconsistencies depends on the existence of appropriate disjointness statements in the ontology the mappings point to. Alternatively, the use of disjointness mappings can provide the same effect. If we want to make clear that chemicals in Tambis are not classified according to the functional view (which we just found to be not entirely true)

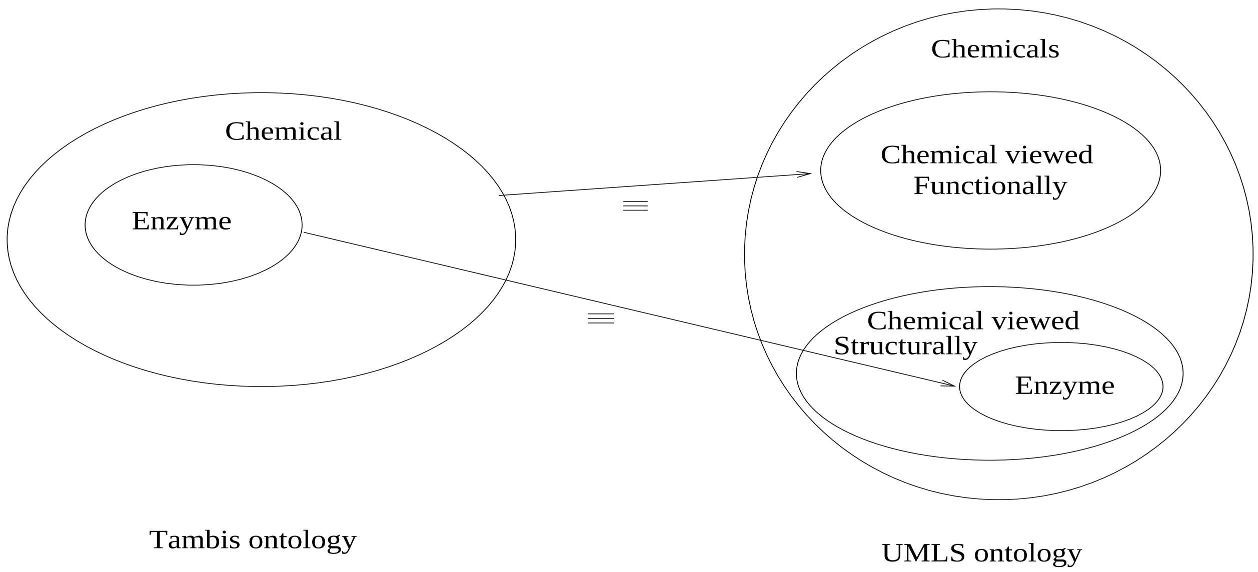

Fig. 5. Inconsistent Mapping in the Example

we can also add a corresponding mapping stating that the image of chemicals is disjoint from the extension of the functional view on chemicals. The definition of this mapping will have the same effect leading to an inconsistency as described above.

Derivation of Mappings Besides the possibility to detect inconsistencies in the mappings,we can also infer additional bridge rules between the same models based on existing ones thereby making the complete mapping implied by the defined rules explicit. We illustrate this possibility by discussing possible implications of an equivalence mapping. Figure 6 illustrates parts of the alignment of substance related alignment of UMLS and GALEN. In particular, it shows the rule stating an equivalence between the GALEN class chemical and the UMLS class chemical substance which is part of the alignment. The definitions in UMLS state that chemical substances are less general than the class generalized substance, more general than complex chemicals and disjoint from processes. As the existing bridge rule states that the image of chemical is exactly the extension of chemical substance in UMLS, these relations also hold between this image and the other UMLS classes mentioned. The relations can be explicated by adding corresponding bridge rules stating that the image of chemicals is more general than complex chemicals, less general that generalized substance and disjoint from processes.

Similar inferences can be made based on bridge rules indicating specialization and generalization relations. If we replace the equivalence in figure 6 by a rule stating that chemicals is more specific than chemical substances, we are still able to infer the relations to generalized substances and to processes. Just the one to complex chemicals will be lost, because the image of chemicals might only overlap or be disjoint from the extension of the respective concept. Conversely, replacing the equivalence by bridge rule stating that chemicals is more general than chemical substances would have preserved the conclusion that chemicals is more general

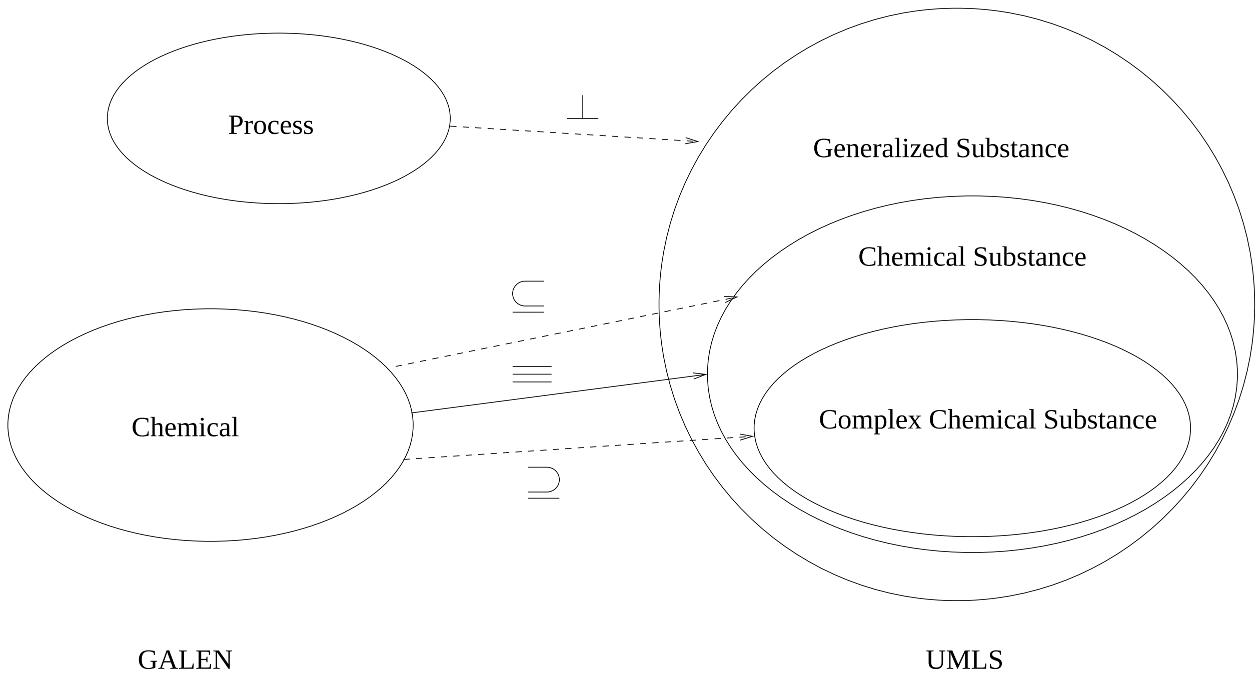

Fig. 6. Derivation of additional Mappings

than complex chemicals. Finally, stating that chemicals is disjoint from chemical substances would have implied that it is also disjoint from complex chemicals.

Merging Local Models Another thing we would like to do based on the alignments is to compare the the local models (Tambis and GALEN) with each other and derive semantic correspondences between classes in these models as well. It turns out that we cannot really drive mappings between the two local models from their mappings to UMLS, because referring to different interpretation domains, we cannot compare the constraints imposed by these mappings. This situation changes, however, when we assume that the local models are to be merged. In this case their interpretation domain becomes the same and we can use the constraints to derive semantic correspondences between concepts in the two models from the existing mappings.

Figure 7 shows two examples of derived relations between concepts from GALEN and Tambis. The figure shows two concepts from each, UMLS (upper part), Tambis (lower left part) and GALEN (lower right part). We assume that we have fixed the inconsistency detected in the mapping from Tambis to UMLS by removing the bridge rule relating chemical substances to the structural view on chemicals and replacing it by an equivalence between chemical substance and chemicals in general. As the GALEN concept chemical is also defined to be equivalent to Chemical, we can derive that these two concepts are equivalent in the merged ontology. Further, we defined the notion of substance in Tambis to be more specific than the same notion in UMLS which is again defined to be more specific than generalized substance in GALEN. From these mappings, we can derive that the Tambis notion of sub-

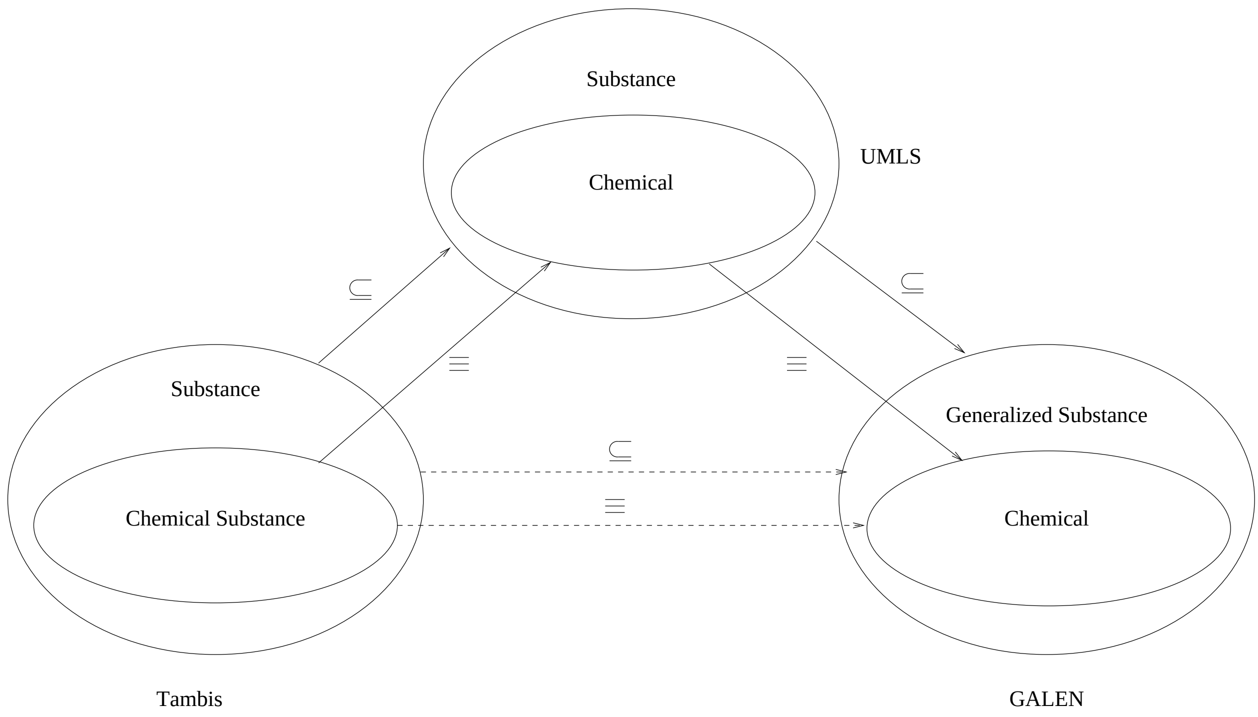

Fig. 7. Derivation of semantic relations in the merged model

stance is more specific than Generalized substance and add a corresponding axiom to the merged ontology.

8 Conclusion

In this paper we have shown how the syntax and the semantics of OWL can be extended to deal with some problems that couldn't otherwise be dealt with. The result is C-OWL (Context OWL), an extended language with an enriched semantics which allows us to contextualize ontologies, namely, to localize their contents (and, therefore, to make them not visible to the outside) and to allow for explicit mappings (bridge rules) which allow for limited and totally controlled forms of global visibility.

This is only the first step and a lot of research remains to be done. The core issue at stake here is the tension between how much we should share and globalize (via ontologies) and how much we should localize with limited and totally controlled forms of globalization (via contexts).

In the last part of this paper we present a first application of C-OWL for the coordination between three complex medical ontologies such as GALEN, Tambis, and UMLS. In this case study it was evident that global sharing the ontologies is inappropriate, as such ontologies are already well established and widely used and sharing would have implied changing them. So we use C-OWL to state semantic mappings between them. Furthermore, we show how, by means of logical reasoning based on C-OWL semantics, additional semantic mappings can be derived on the basis of a set of initial mappings.

References

1. P.G. Baker, C.A. Goble, S. Bechhofer, N.W. Paton, R. Stevens, and A. Brass. An ontology for bioinformatics applications. *Bioinformatics*, 15(6):510–520, 1999.
2. Olivier Bodenreider. The unified medical language system (umls): integrating biomedical terminology. *Nucleic Acids Research*, 32, 2004.
3. M. Bonifacio, P. Bouquet, G. Mameli, and M. Nori. Kex: a peer-to-peer solution for distributed knowledge management. In D. Karagiannis and U. Reimer, editors, *Fourth International Conference on Practical Aspects of Knowledge Management (PAKM-2002)*, Vienna (Austria), 2002.
4. M. Bonifacio, A. Doná, A. Molani, and L. Serafini. Context matching for electronic marketplaces - a case study. *The Knowledge Engineering Review*, 2004. To appear. Short version in Proceedings of the Workshop on Ontologies and Distributed Systems. 18th Int. Joint Conf. on Artificial Intelligence, CEUR Workshop Proceeding series `http://ceur-ws.org/Vol-71/>http://ceur-ws.org/Vol-71`.
5. A. Borgida and L. Serafini. Distributed description logics: Assimilating information from peer sources. *Journal of Data Semantics*, 1:153–184, 2003. Editor in Chief S. Spaccapietra. LNCS 2800, Springer Verlag.
6. P. Bouquet, A. Donà, L. Serafini, and S. Zanobini. ConTeXtualizedlocal ontology specification via ctxml. In P. Bouquet, editor, *Working Notes of the AAAI-02 workshop on Meaning Negotiation. Edmonton (Canada). July 28, 2002*. AAAI, AAAI Press, 2002.
7. P. Bouquet, L. Serafini, and S.Zanobini. Semantic coordination: a new approach and an application. In *Sencond Internatinal Semantic Web Conference*, volume 2870 of *Lecture Notes in Computer Science*, pages 130–145. Springer Verlag, September 2003.
8. J. Davies, D. Fensel, and F. van Harmelen, editors. *Towards the Semantic Web: Ontology-Driven Knowledge Management*. John Wiley, 2002.
9. Aldo Gangemi, Domenico M. Pisanelli, and Geri Steve. Ontology integration: Experiences with medical terminologies. In N. Guarino, editor, *Proceedings of Conference on formal ontologies in information systems (FOIS'98)*. IOS Press, 1998.
10. Gene ontologytm consortium. Web site. `http://www.geneontology.org/`.
11. M. R. Genesereth, A. M. Keller, and O. Duschka. Infomaster: An information integration system. In *Proceedings of 1997 ACM SIGMOD Conference*, 1997.
12. C. Ghidini and F. Giunchiglia. Local models semantics, or contextual reasoning = locality + compatibility. *Artificial Intelligence*, 127(2):221–259, April 2001.
13. C. Ghidini and L. Serafini. Distributed First Order Logics. In D. Gabbay and M. de Rijke, editors, *Frontiers Of Combining Systems 2*, Studies in Logic and Computation, pages 121–140. Research Studies Press, 1998.
14. F. Giunchiglia. Contextual reasoning. *Epistemologia, special issue on I Linguaggi e le Macchine*, XVI:345–364, 1993. Short version in Proceedings IJCAI'93 Workshop on Using Knowledge in its Context, Chambery, France, 1993, pp. 39–49. Also IRST-Technical Report 9211-20, IRST, Trento, Italy.
15. F. Giunchiglia and L. Serafini. Multilanguage first order theories of propositional attitudes. In *Proceedings 3rd Scandinavian Conference on Artificial Intelligence*, pages 228–240, Roskilde University, Denmark, 1991. IOS Press.
16. C. Golbreich, O. Dameron, B. Gibaud, and A. Burgun. Web ontology language requirements w.r.t expressiveness of taxonomy and axioms in medicine. In *The SemanticWeb - ISWC 2003*, volume 2870 of *Lecture Notes in Computer Science*, pages 180–194. Springer Verlag, 2003.
17. D.L. McGuinness and F. van Harmelen. Owl web ontology language overview. `http://www.w3.org/TR/2003/WD-owl-features-20030331/`, March 2003.
18. Stuart J. Nelson and Betsy L. Powell, Tammy andn Humphreys. The unified medical language system (umls) project. In Allen Kent and Carolyn M. Hall, editors, *Encyclopedia of Library and Information Science.*, pages 369–378. Marcel Dekker, Inc., 2002.
19. P.F. Patel-Schneider, P. Hayes, and I. Horrocks. Web Ontology Language (OWL) Abstract Syntax and Semantics. Technical report, W3C, `www.w3.org/TR/owl-semantics/`, February 2003.
20. A.L. Rector and W.A. Nowlan. The galen project. *Computer Methods and Programs in Biomedicine*, 45:75–78, 1993.
21. L. Serafini and F. Giunchiglia. Ml systems: A proof theory for contexts. *Journal of Logic, Language and Information*, 11(4):471–518, 2002.
22. L. Serafini, F. Giunchiglia, J. Mylopoulos, and P. Bernstein. Local relational model: a logical formalization of database coordination. In P. Blackburn, C. Ghidini, and R. Turner, editors, *Proc. of Fourth International and Interdisciplinary Conference on Modeling and Using Context (Context'03)*, LNAI. Springer Verlag, 2003. To appear.