



ELSEVIER

Available online at www.sciencedirect.com

Data & Knowledge Engineering xxx (2007) xxx-xxx

**DATA &
KNOWLEDGE
ENGINEERING**www.elsevier.com/locate/datak

Reasoning and change management in modular ontologies

Heiner Stuckenschmidt ^{a,*}, Michel Klein ^b^a *Universität Mannheim, Germany*^b *Vrije Universiteit Amsterdam, The Netherlands*

Received 13 December 2006; accepted 5 February 2007

Abstract

The benefits of modular representations are well known from many areas of computer science. While in software engineering modularization is mainly a vehicle for supporting distributed development and re-use, in knowledge representation, the main goal of modularization is efficiency of reasoning. In this paper, we concentrate on the benefits of modularization in the context of ontologies, explicit representations of the terminology used in a domain. We define a formal representation for modular ontologies based on the notion of Distributed Description Logics and introduce an architecture that supports local reasoning by compiling implied axioms. We further address the problem of guaranteeing the correctness and completeness of compiled knowledge in the presence of changes in different modules. We propose a heuristic for analyzing changes and their impact on compiled knowledge and guiding the process of updating compiled information that can often reduce the effort of maintaining a modular ontology by avoiding unnecessary re-compilation.

© 2007 Published by Elsevier B.V.

Keywords: Ontologies; Reasoning; Distributed Knowledge Representation; Change management

1. Motivation

Currently, research in the area of the semantic web is in a state where ontologies are ready to be applied in real applications such as semantic web portals, information retrieval or information integration. In order to lower the effort of building ontology-based applications, there is a clear need for a representational and computational infrastructure in terms of general purpose tools for building, storing and accessing ontologies. A number of such tools have been developed, i.e. ontology editors [1,2], reasoning systems [3,4] and more recently storage and query systems (e.g. [5]). Most of these tools, however, treat ontologies as mono-

* Corresponding author. Address: Universität, Mannheim, Institut für Informatik, A5, 6, 68159 Mannheim, DE, Germany. Tel.: +49 621 181 2530.

E-mail address: heiner@informatik.uni-mannheim.de (H. Stuckenschmidt).

28 lithic entities and provide little support for specifying, storing and accessing ontologies in a modular
29 manner.

30 1.1. Why modularization?

31 There are many reasons for thinking about ontology modularization. Our work is mainly driven by three
32 arguments. These also bias the solution we propose, as it focuses on the following aspects.

33 *Distributed Systems:* In distributed environments like the semantic web, the question for modularization
34 arises naturally. Ontologies in different places are built independent of each other and can be assumed
35 to be highly heterogeneous. Unrestricted referencing of concepts in a remote ontology can therefore lead
36 to serious semantic problems as the domain of interpretation may differ even if concepts appear to be
37 the same on a conceptual level. The introduction of modules with local semantics can help to overcome this
38 problem.

39 *Large Ontologies:* Modularization is not only desirable in distributed environments, it also helps to manage
40 very large ontologies that we find in medicine or biology. These ontologies, which sometimes contain more
41 than a hundred thousand concepts, are hard to maintain as changes are not contained locally but can affect
42 large parts of the model. Another argument for modularization in the presence of large ontologies is re-use:
43 in most cases, we are not interested in the complete ontology when building a new system, but only in a
44 specific part. Experiences from software engineering shows that modules provide a good level of abstraction
45 to support maintenance and re-use.

46 *Efficient reasoning:* A specific problem with distributed ontologies as well as with very large models is the
47 efficiency of reasoning. While the pure size of the ontologies causes problems in the latter case, hidden
48 dependencies and cyclic references can cause serious problems in a distributed setting. The introduction
49 of modules with local semantics and clear interfaces will help to analyze distributed systems and provides
50 a basis for the development of methods for localizing inference.

51

52 1.2. Requirements

53 There are three requirements a modular ontology architecture has to fulfill in order to improve ontology
54 maintenance and reasoning in the way suggested above. The requirements will be the main guidelines for
55 the design of our solution proposed in this work.

56 *Loose Coupling:* In general, we cannot assume that two ontology modules have anything in common. This
57 holds for the conceptualization as well as for the interpretation of objects, concepts or relations. Our archi-
58 tecture has to reflect this by providing an extremely loose coupling of modules. In particular, we have to
59 prevent unwanted interactions between modules. For this purpose, mappings between modules have to
60 be distinguished from local definitions on the semantic as well as the conceptual level.

61 *Self-Containment:* In order to facilitate the re-use of individual modules from a larger, possibly intercon-
62 nected system, we have to make sure that modules are self-contained. In particular, it should be possible
63 to perform certain reasoning tasks such as subsumption or query answering within a single module with-
64 out having to access other modules. This is also important if we want to provide efficient reasoning. Fur-
65 ther, we have to ensure correctness, and whenever possible completeness, of local reasoning, for obvious
66 reasons.

67 *Integrity:* The advantages of having self-contained ontology modules have their price in terms of potential
68 inconsistencies that arise from changes in other ontology modules. While there is in our architecture no
69 need to access other modules at reasoning time, the correctness of reasoning within a self contained module
70 may still depend on knowledge in other ontologies. If this knowledge changes, reasoning results in a self-
71 contained module may become incorrect with respect to the overall system, and we will not even notice it.

72 We have to provide mechanisms for checking whether relevant knowledge in other systems has changed
73 and for adapting the reasoning process if needed to ensure correctness.
74

75 1.3. Related work

76 Our work relates to two main areas of research on representing and reasoning about ontological knowl-
77 edge. The first is concerned with distributed and modular knowledge representation where we use ideas from
78 theorem proving and knowledge engineering. The second area of related work is concerned with managing
79 knowledge models. Here previous work exists in the area of knowledge engineering and semantic web
80 technologies.

81 While the principle of modularity has widely been adopted in software engineering it has got less attention
82 in the area of knowledge representation and reasoning. Some fundamental work on the modularization of rep-
83 resentations can be found in the area of theorem proving. Farmer and colleagues promote the use of combi-
84 nations of ‘Little Theories’, representations of a specific mathematical structure in order to reason about
85 complex problems [6]. They show the advantages of this modular approach in terms of reusability and reduced
86 modeling effort. The idea of reusing and combining chunks of knowledge rather than building knowledge
87 bases from scratch has later been adopted by the knowledge engineering community for building real-world
88 knowledge bases (e.g. see [7]). McIlraith and Amir argue that a modularization of knowledge bases has also
89 advantages for reasoning, even if the modularization is done a posteriori. They present algorithms for break-
90 ing down existing representations into a set of modules with minimal interaction and define reasoning proce-
91 dures for propositional [8] and first-order logic [9]. The work reported there is motivated by well established
92 techniques from uncertain reasoning, where an a posteriori modularization of large theories is a common way
93 to reduce runtime complexity (e.g. see [10]).

94 As we are interested in representations of ontological knowledge, approaches from the area of logics for
95 representing terminologies, so-called description logics are of special interest for our work. In this area, we
96 find the same arguments for a modularized representation as in the area of theorem proving. Rector pro-
97 poses a strategy for modular implementation of ontologies using description logics [11]. The approach is
98 based on a set of orthogonal taxonomies that provide a basis for defining more complex concepts. Rector
99 argues for the benefits of this strategy in terms of easier creation and re-use of ontological knowledge. Buch-
100 heit and others propose a similar structuring on the language level by dividing the terminological part of a
101 knowledge base into a schema part that corresponds to the basic taxonomies and a view part [12]. They
102 show that this distinction can be used to achieve better run-time behavior for complex view languages.
103 While these approaches still assume the overall model to be a single ontology providing a coherent concep-
104 tualization of the world, Giunchiglia and others propose a more radical approach to distributed represen-
105 tations. They propose the local model semantics as an extension of the standard semantics of first order
106 logics [13]. This semantics allows different modules to represent different views on the same part of the
107 world and the definition of directed partial mappings between different modules. Recently, Borgida and Ser-
108 afini defined a distributed version of description logics based on local model semantics that has all advan-
109 tages of the contextual representations [14].

110 As already mentioned, the problem of combining and reasoning with ontological modules has become of
111 central importance in research on knowledge representation and reasoning on the semantic web. Standard lan-
112 guages for encoding ontological knowledge on the World Wide Web, i.e. the RDF schema [15] and the web
113 ontology language OWL [16] provide some basic mechanisms for combining modular representations. The
114 abilities to combine different models are restricted to the import of complete models and to the use of elements
115 from a different model in definitions by direct reference. It is assumed that references to external statements are
116 only made for statements from imported models, however, this is strictly speaking not required. As a conse-
117 quence, mappings rather implicitly exist in terms of mutual use of statements across models. In [17] the
118 authors provide a detailed analysis of the drawbacks of using OWL import statements for combining hetero-
119 geneous ontologies and show that local model semantics, which is also the basis for our work solves these
120 problems. Volz and colleagues discuss different interpretations of the import statement that range from purely
121 syntactic to schema-aware interpretations of the imported knowledge [18]. An alternative way of relating dif-

ferent RDF models to each others that is much closer to our ideas is discussed by Oberle [19] who defines a view language for RDF.

Recently, there has been some interest in formal models for modular ontologies that can be seen as competing approaches to the one described in this paper. In particular, researchers have proposed the use of ϵ -connections [20] as a suitable formalism for inter-module links and investigated the logical properties of ontologies. Others have proposed to strengthen distributed description logics by adding requirements to the links between different local models resulting in a formalism called P-DL [21]. These models, however, take a slightly different view on modular ontologies. While in this paper, we assume that modular ontologies are created by linking previously unrelated, possibly inconsistent ontologies that can also have overlap in their scope, the above mentioned approaches focus on a scenario where an existing ontology is partitioned into a number of modules for the sake of enhancing (re-)usability and thus assume a tighter coupling of the different modules which requires a stronger formalism for specifying links between the modules.

1.4. Our approach

In the following, we describe our approach to ontology modularization on an abstract level. We emphasize the main design decisions and motivate them on the basis of the requirements defined above. The technical details of the approach will be given in the subsequent sections.

View-Based Mappings: The first design decision concerns the way different ontology modules are connected. In our work, we adopt the approach of view-based information integration. In particular, ontology modules are connected by conjunctive queries and the extension of a concept in one module can be claimed to be equivalent to the (intentional) answer set of a conjunctive query over the vocabulary of another module. This way of connecting modules is more expressive than simple one-to-one mappings between concept names. Further, the same technique can be used to define relations of any arity based on other modules. Compared to the use of arbitrary axioms, our approach is less expressive. We decide to sacrifice a higher expressiveness for the sake of conceptual simplicity and desirable semantic properties that are discussed in the remainder of this paper.

Interface Compilation: The use of conjunctive queries guarantees a loose coupling on a conceptual and semantic level. However, it does not provide self-containment, because reasoning in an ontology module depends on the answer sets of the queries that are used to connect it to other modules. These answer sets have to be determined by actually querying the other ontology module. In order to make local reasoning independent from other modules, we use a knowledge compilation approach. The idea is to compute the result of each mapping query off-line and add the result as an axiom to the ontology module. At reasoning time these axioms replace the query, thus enabling local reasoning. As the results of queries are considered to be defined intentionally rather than extensionally, the result of the compilation of a query is not a set of instances retrieved from other modules, but a set of axioms that contains all the information necessary to perform local reasoning.

Change Detection and Automatic Update: Our approach of compiling mappings and adding the result to the different ontology modules is very sensitive against changes in ontology modules. Once a query has been compiled, the correctness of reasoning can only be guaranteed as long as the queried ontology module does not change. On the other hand, not every change in the system does really influence the compiled result. Problems only arise if concepts used in the query change or if the set of concepts subsuming the query is changed. In the second case, we will have to compile the interface again. In the first case we might even have to consider a redefinition of the query. In order to decide whether the compiled axiom is still valid, we propose a change detection mechanism that is based on a taxonomy of ontological changes and their impact on the concept hierarchy in combination with the position of the affected concept in the hierarchy.

In the following, we first introduce a representational framework for modular ontologies that builds on top of existing work on distributed description logics (DDL) as a framework for reasoning about distributed ontologies. In Section 3 we define reasoning mechanisms for modular ontologies as a special case of general inference in distributed description logics. We further introduce the compilation of implied subsumption rela-

Table 1
Axiom patterns for describing ontologies in description logics

DL Axiom	Semantics	Intuition
A-Box		
$C(x)$	$x^{\mathcal{I}} \in C^{\mathcal{I}}$	x is of type C
$P(x, y)$	$(x^{\mathcal{I}}, y^{\mathcal{I}}) \in P^{\mathcal{I}}$	x is related to y by R
T-Box		
$C \sqsubseteq D$	$C^{\mathcal{I}} \subseteq D^{\mathcal{I}}$	C is more specific than D
$p \sqsubseteq r$	$p^{\mathcal{I}} \subseteq r^{\mathcal{I}}$	P is more specific than R
$p \equiv r^{-}$	$p^{\mathcal{I}} = \{(x, y) (y, x) \in r^{\mathcal{I}}\}$	P is the inverse of R

171 tions as a mechanisms for localizing reasoning and compare it with the distributed reasoning methods pro-
 172 posed for DDL. Section 4 discusses the problem of handling changes in external ontologies and their impact
 173 on compiled knowledge and proposes a heuristic for checking whether compiled knowledge has to be recom-
 174 puted. We conclude with an example from a case study on ontology evolution in Section 5 and a discussion of
 175 the tradeoffs of our approach and possible extensions in Section 7.

176 2. Modular ontologies

177 In this section, we present a formal model for modular ontologies that will be used throughout the paper.
 178 Our starting point is the use of description logics – a special kind of logics for representing terminological
 179 knowledge as the basis for representing ontologies. We briefly explain the nature of description logics and their
 180 semantics. We then formally introduce the logic \mathcal{SHIQ} which is the basis for our work. We then proceed
 181 with the definition of our model for modular ontologies by briefly recalling an extension of \mathcal{SHIQ} with map-
 182 pings between different models known as Distributed Description Logics (DDL). As our model turns out to be
 183 a subset of Distributed Description Logics, we conclude this section by explaining the restrictions to the gen-
 184 eral framework of DDL that apply to the modular ontologies.

185 2.1. Ontologies and description logics

186 An Ontology usually groups objects of the World that have certain properties in common (e.g. cities or
 187 countries) into concepts. A specification of the shared properties that are characteristic for this set of objects
 188 is called a concept definition. Concepts can be arranged into a subclass–superclass relation in order to be able
 189 to further discriminate objects into subgroups (e.g. capitals or European countries). Concepts can be defined
 190 in two ways, by enumeration of its members or by a concept expression. The specific logical operators that can
 191 be used to formulate concept expressions can vary between ontology languages (Table 1).

192 Description Logics are decidable subsets of first order logic that are designed to describe concepts in terms
 193 of complex logical expressions¹ The basic modeling elements in Description Logics are concepts (classes of
 194 objects), roles (binary relations between objects) and individuals (named objects). Based on these modeling
 195 elements, Description Logics contain operators for specifying so-called concept expressions that can be used
 196 to specify necessary and sufficient conditions for membership in the concept they describe. These modeling
 197 elements are provided with a formal semantics in terms of an abstract domain interpretation mapping \mathcal{I} map-
 198 ping each instance onto an element of an abstract domain $\Delta^{\mathcal{I}}$. Instances can be connected by binary relations
 199 defined as subsets of $\Delta^{\mathcal{I}} \times \Delta^{\mathcal{I}}$. Concepts are interpreted as a subset of the abstract domain Δ . Intuitively, a
 200 concept is a set of instances that share certain properties. These properties are defined in terms of concept
 201 expressions. Typical operators are the Boolean operators as well as universal and existential quantification
 202 over relations to instances in other concepts.

203 A Description Logic Knowledge base consists of two parts. The A-box contains information about objects,
 204 their type and relations between them, the so-called T-Box consists of a set of axioms about concepts (poten-
 205 tially defined in terms of complex concept expressions and relations). The first type of axioms can be used to

¹ Details about the relation between description logics and first order logic can be found in [22,23].

206 describe instances. In particular, axioms can be used to state that an instance belongs to a concept or that two
 207 instances are in a certain relation. It is easy to see, that these axioms can be used to capture case descriptions as
 208 labeled graphs. The other types of axioms describe relations between concepts and instances. It can be stated
 209 that one concept is a subconcept of the other (all its instances are also instances of this other concept). Further,
 210 we can define a relation to be a subrelation or the inverse of another relation. These axioms are used to for-
 211 malize the legal ontology described in the last section.

212 The formal semantics of concepts and relations as defined by the interpretation into the abstract domain $\Delta^{\mathcal{I}}$
 213 can be used to automatically infer new axioms from existing definitions.

214 It has been argued that encoding ontologies in Description Logics is beneficial, because it enables inference
 215 engines to reason about ontological definitions. In this context, deciding subsumption between two concept
 216 expressions, i.e. deciding whether one expression is more general than the other one is one of the most impor-
 217 tant reasoning tasks as it has been used to support various tasks including information integration [24], prod-
 218 uct and service matching [25,26] and Query answering over ontologies [27].

219 In this paper, we consider ontologies represented in the description logic \mathcal{SHIQ} . This choice is moti-
 220 vated by the fact that \mathcal{SHIQ} covers a large part of the expressive power of the Web Ontology Language
 221 OWL [16], more specifically of the language OWL-DL, a decidable sublanguage of OWL that directly cor-
 222 responds to the logic \mathcal{SHOIQ} that extends \mathcal{SHIQ} with nominals [28]. We omit this extension in order to
 223 be able to base our framework on recent results on Distributed Description Logics [29,30] that provide us
 224 with basic mechanisms for specifying links between concepts in different ontologies in a loose way. Before
 225 defining our notion of modular ontologies, we briefly introduce the logic \mathcal{SHIQ} as well as the basic
 226 notions of Distributed Description Logics. For further information about notation and naming in Descrip-
 227 tion Logics, we refer to [31].

228 2.2. The \mathcal{SHIQ} description logic

229 Let \mathcal{C} be a set of concept names and RN a set of role names. Further let there be a set $R^+ \subseteq RN$ of transitive
 230 roles (i.e. for each $r \in R^+$ we have $r(x, y) \wedge r(y, z) \Rightarrow r(x, z)$). If now r^- denotes the inverse of a role (i.e.
 231 $r(x, y) \Rightarrow r^-(y, x)$) then we define the set of roles R as $RN \cup \{r^- \mid r \in RN\}$. A role inclusion axiom is an expres-
 232 sion $r \sqsubseteq s$ where r and s are roles. A role is called a simple role if it is not transitive and does not have transitive
 233 subroles with respect to the transitive closure of the role inclusion relation. Concept expressions are now
 234 formed by applying special operators to concept and role names. In particular, new concept expressions
 235 can be formed from existing ones using the Boolean operators or by imposing constraints on the type and
 236 number of objects related to objects of the described concept. The corresponding operators are summarized
 237 below

239 Expression	Intuition
242 $\neg C$	All objects that are not of type C
243 $C \sqcap D$	All objects that are of type C and of type D
246 $C \sqcup D$	All Objects that are of type C or of type D
248 $\exists r.C$	All Objects that related to some objects of type C via 249 relation r
250 $\forall r.C$	All Objects that are only related to objects of type C via 252 relation r
253 $(\geq nr.C)$	All objects that are related to at least n objects of type C 255 via relation r
256 $(\leq nr.C)$	All objects that are related to at most n objects of type C 258 via relation r

259

260 Formally, the set of concepts (or concept expressions) in \mathcal{SHIQ} is the smallest set such that:

- 261 • \top and \perp are concept expressions for the most general concept and the unsatisfiable concept, respectively;

- 262 • every concept name A is a concept expression;
 263 • if C and D are concept expressions, r is a role, s is a simple role and n is a non-negative integer, then $\neg C$,
 264 $C \sqcap D$, $C \sqcup D$, $\forall r.C$, $\exists r.C$, $(\geq nr.C)$ and $(\leq nr.C)$ are concept expressions.

265

266 A general concept inclusion axiom is an expression $C \sqsubseteq D$ where C and D are concepts. A terminology is a
 267 set of general concept inclusion and role inclusion axioms.

268 The semantics of \mathcal{SHIQ} is defined in terms of an interpretation $\mathcal{I} = (\Delta^{\mathcal{I}}, \cdot^{\mathcal{I}})$ where $\cdot^{\mathcal{I}}$ is a function
 269 that maps every concept on a subset of $\Delta^{\mathcal{I}}$ and every role on a subset of $\Delta^{\mathcal{I}} \times \Delta^{\mathcal{I}}$ such that for all concepts
 270 C and D and for roles r where $\#M$ denotes the cardinality of M and $(r^{\mathcal{I}})^+$ the transitive closure of $r^{\mathcal{I}}$ we
 271 have

- 272 • $\top^{\mathcal{I}} = \Delta^{\mathcal{I}}$ and $\perp^{\mathcal{I}} = \emptyset$
 273 • $r^{\mathcal{I}} = (r^{\mathcal{I}})^+$ for $r \in R^+$ and $r^- = \{(y, x) \mid (x, y) \in r^{\mathcal{I}}\}$
 274 • $(\neg C)^{\mathcal{I}} = \Delta^{\mathcal{I}} - C^{\mathcal{I}}$, $(C \sqcap D)^{\mathcal{I}} = C^{\mathcal{I}} \cap D^{\mathcal{I}}$ and $(C \sqcup D)^{\mathcal{I}} = C^{\mathcal{I}} \cup D^{\mathcal{I}}$
 275 • $(\forall r.C)^{\mathcal{I}} = \{x \mid \forall y. (x, y) \in r^{\mathcal{I}} \Rightarrow y \in C^{\mathcal{I}}\}$
 276 • $(\exists r.C)^{\mathcal{I}} = \{x \mid \exists y. (x, y) \in r^{\mathcal{I}} \wedge y \in C^{\mathcal{I}}\}$
 277 • $(\geq nr.C)^{\mathcal{I}} = \{x \mid \#\{y. (x, y) \in r^{\mathcal{I}} \wedge y \in C^{\mathcal{I}}\} \geq n\}$
 278 • $(\leq nr.C)^{\mathcal{I}} = \{x \mid \#\{y. (x, y) \in r^{\mathcal{I}} \wedge y \in C^{\mathcal{I}}\} \leq n\}$

279

280 An interpretation satisfies a terminology \mathcal{T} if $C^{\mathcal{I}} \subseteq D^{\mathcal{I}}$ for all general concept inclusions $C \sqsubseteq D$ in \mathcal{T} and
 281 $r^{\mathcal{I}} \subseteq s^{\mathcal{I}}$ for all role inclusion axioms $r \sqsubseteq s$ in \mathcal{T} . In this case we call \mathcal{I} a model for \mathcal{T} . A concept D subsumes a
 282 concept C in \mathcal{T} if $C \sqsubseteq D$ holds for all models of \mathcal{T} . In the remainder of the paper we will focus on the task of
 283 deciding whether a concept subsumes another one.

284 2.3. Distributed description logic

285 Distributed Description Logics as proposed in [29] provide a language for representing sets of terminologies
 286 and semantic relations between them. For this purpose DDLs provide mechanisms for referring to terminol-
 287 ogies and for defining rules that connect concepts in different terminologies. On the semantic level, DDLs
 288 extend the notion of interpretation introduced above to fit the distributed nature of the model and to reason
 289 about concept subsumption across terminologies.

290 Let I be a non-empty set of indices and $\{\mathcal{T}_i\}_{i \in I}$ a set of terminologies. We prefix inclusion axioms with the
 291 index of the terminology they belong to (i.e. $i : C$ denotes a concept in terminology \mathcal{T}_i and $j : C \sqsubseteq D$ a concept
 292 inclusion axiom from terminology \mathcal{T}_j). Note that $i : C$ and $j : C$ are different concepts. Semantic relations
 293 between concepts in different terminologies are represented in terms of axioms of the following form, where
 294 C and D are concepts in terminologies \mathcal{T}_i and \mathcal{T}_j , respectively:

- 295 • $i : C \stackrel{\sqsubseteq}{\rightarrow} j : D$ (into-rule)
 296 • $i : C \stackrel{\supseteq}{\rightarrow} j : D$ (onto-rule)

297

298 These axioms are also called *bridge-rules*. The into-rule states that concept C in terminology \mathcal{T}_i is intended
 299 to be more specific than concept D in terminology \mathcal{T}_j . Conversely, the onto-rule states that concept C in ter-
 300 minology \mathcal{T}_i is intended to be more general than concept D in terminology \mathcal{T}_j . An additional rule $i : C \stackrel{\equiv}{\rightarrow} j : D$
 301 is defined as the conjunction of the two rules above, stating that the two concepts are intended to be equiv-
 302 alent. A distributed terminology \mathfrak{T} is now defined as a pair $(\{\mathcal{T}_i\}_{i \in I}, \{\mathfrak{B}_{ij}\}_{i \neq j \in I})$ where $\{\mathcal{T}_i\}_{i \in I}$ is a set of ter-
 303 minologies and $\{\mathfrak{B}_{ij}\}_{i \neq j \in I}$ is a set of bridge rules between these terminologies.

304 The semantics of distributed description logics is defined in terms of a global interpretation
 305 $\mathfrak{I} = (\{\mathcal{I}_i\}_{i \in I}, \{r_{ij}\}_{i \neq j \in I})$ where \mathcal{I}_i is an interpretation for terminology \mathcal{T}_i as defined above and
 306 $r_{ij} \subseteq \Delta^{\mathcal{I}_i} \times \Delta^{\mathcal{I}_j}$ is a domain relation connecting elements of the interpretation domains of terminologies \mathcal{T}_i
 307 and \mathcal{T}_j . We use $r_{ij}(x)$ to denote $\{y \in \Delta^{\mathcal{I}_j} \mid (x, y) \in r_{ij}\}$ and $r_{ij}(C)$ to denote $\bigcup_{x \in C} r_{ij}(x)$.

308 A distributed interpretation \mathfrak{I} satisfies a distributed terminology \mathfrak{T} if:

- 309 • \mathcal{I}_i satisfies \mathcal{T}_i for all $i \in I$
 310 • $r_{ij}(C^{\mathcal{I}_i}) \subseteq D^{\mathcal{I}_j}$ for all $i : C \xrightarrow{\Xi} j : D$ in \mathfrak{B}_{ij}
 311 • $r_{ij}(C^{\mathcal{I}_i}) \supseteq D^{\mathcal{I}_j}$ for all $i : C \xrightarrow{\exists} j : D$ in \mathfrak{B}_{ij}
 312

313 In this case we call \mathfrak{T} a model for \mathfrak{T} . A concept $i : D$ subsumes a concept $i : C$ ($i : C \sqsubseteq D$) if for all models of
 314 \mathfrak{T} we have $C^{\mathcal{I}_i} \subseteq D^{\mathcal{I}_i}$.

315 2.4. Modular ontologies

316 We can now define our notion of a modular ontology in terms of Distributed Description Logics. In fact,
 317 our notion of a modular ontology is a restricted form of distributed terminology as defined above. The restric-
 318 tions we impose concern the architecture of the distributed terminology as well as the expressiveness of seman-
 319 tic relations between terminologies. These restrictions are motivated by the aims of (1) providing an alternative
 320 to the standard notion of import in OWL and (2) the goal of providing support for localized reasoning and
 321 maintenance of the modular ontology. In the following, we will first discuss the architecture of a modular
 322 ontology and then introduce the restrictions imposed on semantic relations.

323 2.4.1. Architecture

324 As described above, DDL makes a clear distinction between terminologies and semantic mappings between
 325 them in terms of bridge rules, which in principle are independent of the terminologies. This makes the model
 326 quite flexible; for example, it permits having different sets of mapping rules connecting the same set of ontol-
 327 ogies. In this way it is possible to encode different views on how the terminologies relate to each other. In con-
 328 trast, our aim is to enable the use of external knowledge in a terminology similar to the ability of OWL to use
 329 concept and role names defined in different terminologies. This view is different from the model of Distributed
 330 Description Logics as it makes the semantic links to other models part of the terminology. Being part of the
 331 terminology implies that there is only one way of connecting to these external definitions which is assumed to
 332 be agreed on by the users of the local terminology.

333 We achieve this localization of semantic relations by introducing the notion of externally defined concepts
 334 in a terminology. We divide the set of concept names in a terminology into internally defined concepts \mathcal{C}_I and
 335 externally defined concepts \mathcal{C}_E resulting into the following description of the set of all concept names
 336 \mathcal{C} ($\mathcal{C} = \mathcal{C}_I \cup \mathcal{C}_E$, $\mathcal{C}_I \cap \mathcal{C}_E = \emptyset$).

337 We consider externally defined concepts to be concept names linked to a concept expression defined in
 338 another terminology using bridge rules. An external concept definition in terminology \mathcal{T}_i is an axiom of the
 339 form: $i : C \equiv \mathcal{T}_j : D$ where C is a concept name in \mathcal{T}_i , \mathcal{T}_j is a terminology different from the one in which
 340 the external concept is defined and D is a concept expression in \mathcal{T}_j . Note that although D is syntactically rep-
 341 resented in \mathcal{T}_i it actually represents a concept in \mathcal{T}_j . In particular the expression D is only allowed to contain
 342 concepts defined in \mathcal{T}_j . This definition is very close to the OWL mechanism of using concept and role names
 343 from other name spaces in definitions.

344 We give external concept definitions a semantics in terms of distributed description logics by defining exter-
 345 nal concept definitions to be an alternative notation for a pair of bridge rules:

$$347 \quad i : C \equiv \mathcal{T}_j : D \iff j : D \xrightarrow{\Xi} i : C \wedge j : D \xrightarrow{\exists} i : C$$

348 The correspondence between external concept definitions and bridge rules allows us to base our further inves-
 349 tigation on the formal results that have been established for distributed **SHIQ** terminologies.

350 2.4.2. Restricting mapping expressiveness

351 In Distributed Description Logics, there are no restrictions on the antecedent of a bridge rule—except that
 352 it has to be a valid concept of the source terminology. In our framework, we restrict this freedom for the sake
 353 of an easier maintenance of the semantic relations between terminologies. This restriction is motivated by our
 354 earlier work on keeping integrity in modular ontologies reported in [32]. In that work, we proposed a heuristic
 355 approach for determining the impact of changes in other modules on the correctness of local subsumption

356 reasoning. The approach relied on the fact that changes were only monotonically propagated to other mod-
 357 ules. In order to achieve this effect also in the framework of distributed description logics, we restrict the lan-
 358 guage used to specify externally defined concepts to a sublanguage of \mathcal{SHIQ} that does not contain operators
 359 that can have a non-monotonic effect, in particular negation, universal restrictions and qualified number
 360 restrictions that limit the number of related concepts. More precisely, we allow concept expressions that
 361 are defined in the following sublanguage of \mathcal{SHIQ} :

$$363 \quad C, D \rightarrow \top \mid \perp \mid A \mid C \sqcap D \mid C \sqcup D \mid \exists r.C \mid \exists r^-.C \mid \geq nr.C \mid \geq nr^-.C$$

364 In order to restrict the semantic correspondences between terminologies in our model, we now only allow the
 365 concept expressions D in the definition of external concepts to be valid concepts over terminology \mathcal{T}_j with
 366 respect to the sublanguage defined above. We denote such concepts as D_2 and consider external concept
 367 expressions of the form $i : C \equiv j : D_2$.

368 3. Reasoning in modular ontologies

369 The direct correspondence of our framework to Distributed Description Logic allows us to base inference
 370 in modular ontologies on known results for the corresponding DDL. In particular, we can provide complete-
 371 ness and complexity results for reasoning in modular ontologies. We extend the existing work on reasoning in
 372 DDL with the notion of compilation of implied subsumption relations. Specifically, we use reasoning methods
 373 for Distributed Description Logics to derive subsumption relations between externally defined concepts in
 374 modules and explicitly add the derived subsumption relations as axioms to the module. The results of [30]
 375 guarantee that after this compilation step reasoning can be performed locally without considering other mod-
 376 ules unless there are changes in the system.

377 In the following, we first briefly review basic definitions of reasoning in distributed description logics and
 378 prove that it has the same worst-case complexity as reasoning in \mathcal{SHIQ} . We then present the compilation of
 379 subsumption relations and discuss conditions for completeness and consistency.

380 3.1. Reasoning based on DDL

381 Reasoning in DDL differs from reasoning in traditional Description Logics by the way knowledge is prop-
 382 agated between T-Boxes by certain combinations of bridge rules. The simplest case in which knowledge is
 383 propagated is the following:

$$385 \quad \frac{i : A \xrightarrow{\exists} j : G, i : B \xrightarrow{\sqsubseteq} j : H, i : A \sqsubseteq B}{j : G \sqsubseteq H} \quad (1)$$

386 This means that the subsumption between two concepts in a T-Box can depend on the subsumption between
 387 two concepts in a different T-Box if the subsumed concepts are linked by the *onto*- and the subsuming concepts
 388 by an *into* rule. In languages that support disjunction, this basic propagation rule can be generalized to sub-
 389 sumption between a concept and a disjunction of other concepts in the following way:

$$392 \quad \frac{i : A \xrightarrow{\exists} j : G, i : B_k \xrightarrow{\sqsubseteq} j : H_k (1 \leq k \leq n), i : A \sqsubseteq \bigsqcup_{k=1}^n B}{j : G \sqsubseteq \bigsqcup_{k=1}^n H_k} \quad (2)$$

393 It has been shown that this general propagation rule completely describes reasoning in DDLs that goes be-
 394 yond well known methods for reasoning in Description Logics. To be more specific, adding the inference rule
 395 in Eq. 2 to existing tableaux reasoning methods leads to a correct and complete method for reasoning in
 396 DDLs. A corresponding result using a fixpoint operator is given in [30]. Based on these results, we can define
 397 a general inference rule for the case of modular ontologies in the following way:

$$399 \quad \frac{i : A \equiv j : G, i : B_k \equiv j : H_k (1 \leq k \leq n), i : A \sqsubseteq \bigsqcup_{k=1}^n B}{j : G \sqsubseteq \bigsqcup_{k=1}^n H_k} \quad (3)$$

400 There are a number of consequences of this result for reasoning in modular ontologies.

401 *Correctness and Completeness*: From the basic propagation rule, we can see that subsumption between
402 externally defined concepts follows from subsumption of their definitions in the (same) external module. This
403 is because each external concept definition corresponds to an *into* and an *onto* rule between the concept name
404 and its definition. The language we consider is \mathcal{SHIQ} and therefore we have to consider the general prop-
405 agation rule because we have disjunction in our language. This means that it is not enough to simply check
406 whether subsumption between the definitions of two externally defined concepts in the external module is com-
407 plete, but we have to consider all subsets of the set of external concepts. We will discuss this point in more
408 detail in the next section.

409 *Complexity*: As we reduce reasoning in modular ontologies to reasoning in DDLs with \mathcal{SHIQ} as a local
410 language, complexity results can be derived from known results on reasoning in \mathcal{SHIQ} and Distributed
411 Description Logics.

412 **Theorem 1** (Complexity). *Reasoning in modular ontologies is Exp-Time Complete.*

413 **Proof.** We show that reasoning in modular ontologies has the same complexity as reasoning in \mathcal{SHIQ} . As
414 reasoning in \mathcal{SHIQ} is Exp-Time Complete [33], this establishes the result.

- 415 • Reasoning in modular ontologies is at least as hard as reasoning in \mathcal{SHIQ} ; in the extreme case a modular
416 ontology consists only of a single module without external concepts, thus reasoning in modular ontologies
417 is equivalent to reasoning in \mathcal{SHIQ} .
- 418 • Reasoning in modular ontologies is not harder than reasoning in \mathcal{SHIQ} because we reduce reasoning in
419 modular ontologies to reasoning in DDLs. There exists a reduction of reasoning in DDLs with \mathcal{SHIQ} as a
420 local language to \mathcal{SHIQ} [30]. Both reductions are linear in the size of the resulting terminology and there-
421 fore do not change the complexity class.

422

423 This result shows that the complexity of reasoning in modular ontologies is not worse than reasoning in the
424 web ontology language. Using the reduction of DDL to \mathcal{SHIQ} it is even possible to use existing OWL rea-
425 soners for reasoning with modular ontologies. Although practical implementations of OWL reasoners have
426 shown that good average case performance can be achieved, the worst case complexity is still very high
427 and asks for further optimization.

428 3.2. Compilation and integrity

429 Existing reasoners for expressive Description Logics are highly optimized with respect to deciding sub-
430 sumption in the context of a single T-Box. Serafini and Tamilin present a distributed reasoning system that
431 extends existing reasoners to distributed T-Boxes [34]. In theory, this system is complete with respect to the
432 propagation rules described above and has—as we have argued—the same worst-case complexity. In practice,
433 however, reasoning with multiple, possible distributed modules, brings some new problems with respect to
434 completeness and reasoning performance. First of all the completeness of the distributed reasoners depends
435 on the availability of local reasoners for all T-Boxes in the system. In a loosely coupled network without cen-
436 tral control this cannot always be guaranteed as network nodes can be unreachable or even leave the network.
437 In this case, necessary subsumption tests cannot be performed at these nodes leading to a possible incompleteness.
438 Another problem currently not addressed in the work of Serafini and Tamilin are performance problems
439 due to communication costs between the different nodes in the system. Work in the area of distributed dat-
440 abases has shown that communication costs often become serious bottlenecks in distributed systems.

441 In order to overcome these problems we propose to compute subsumption relations between external con-
442 cepts offline and store them as explicit axioms in the local ontologies. If we compute these relations using the
443 reasoner mentioned above we have the guarantee that reasoning about subsumption in each module can be
444 done without caring about the availability of other nodes in the network. This also has the advantage that
445 no communication costs occur as part of online reasoning.

446 Of course these runtime benefits have their price in terms of computational complexity of the compilation
447 step. The completeness of the propagation rule given in Eq. 2 tells us that to be independent from other mod-

ules we only have to consider subsumption relations between externally defined concepts, as only such subsumption relations can be propagated from outside. What we have to check is subsumption between each external concept and the disjunction of all combinations of other external concepts. For a local module, this process is defined in Algorithm 1.

Algorithm 1. Compile

```

453 Require: An T-Box  $\mathcal{T}$  with external concepts  $\mathcal{C}_E$ 
454 forall  $c \in \mathcal{C}_E$  do
455   candidates :=  $\mathcal{P}(\mathcal{C}_E - \{c\})$ 
456   forall  $d \in$  candidates do
457     If  $\mathcal{I} \models c \sqsubseteq \bigsqcup_{e \in d} e$  then
458        $\mathcal{T} := \mathcal{T} \cup \{c \sqsubseteq \bigsqcup_{e \in d} e\}$ 
459     end if
460   end for
461 end for
462
463
464
465

```

If we denote the number of external concepts \mathcal{C}_E as n , the worst-time complexity of the compilation method is $O(n \cdot 2^{(n-1)})$ as can easily be seen from the algorithm. As deciding the subsumption relation in the conditional statement of the algorithm itself is already Exp-Time Complete and this test has to be carried out an exponential number of times with respect to the number of external concepts, compiling all implied statements is computationally very expensive. We therefore do not want to perform the compilation step more often than absolutely necessary to guarantee that local reasoning is still complete.

While the results of Serafini and others [34] guarantee that local reasoning is correct and complete at the time the compilation is carried out, a problem occurs when changes are made to the system. Changes in the definitions of the external concepts, but also changes in the definitions of concepts and relations in other modules can make local reasoning incomplete or inconsistent. In order to prevent situations in which local reasoning is not correct and complete any more we introduce the notion of integrity of a modular ontology.

Definition 1 (Integrity). Let $\mathfrak{I} = (\{\mathcal{T}_i\}_{i \in I}, \{\mathfrak{B}_{ij}\}_{i \neq j \in I})$ be a modular ontology with interpretation $\mathfrak{I} = (\{\mathcal{I}_i\}_{i \in I}, \{r_{ij}\}_{i \neq j \in I})$ then we say that integrity holds for \mathfrak{I} if for all $\mathcal{T}'_i = \text{compile}(\mathcal{T}_i)$ with interpretation \mathcal{I}'_i we have:

$$\mathcal{I}'_i \models C \sqsubseteq D \iff \mathfrak{I} \models i : C \sqsubseteq D$$

for any pair of legal concept expressions C and D in \mathcal{T}'_i .

The notion of integrity gives us a criterion for deciding whether compiled results are still valid. What the definition does not provide is an operational account for checking it. A direct use of the definition would involve a complete check of all derivable subsumption relations. As we have argued above this approach is extremely expensive. In the following, we therefore present a heuristic approach for checking integrity in modular ontologies that is driven by changes made to the ontology. The approach is capable of determining situations in which changes to a modular ontology do not affect integrity and therefore no re-compilation is necessary.

4. Evolution management

As we have argued above, guaranteeing integrity of compiled subsumption relations is the main problem in modular ontologies. In principle, all compiled subsumption relations have to be recomputed to test whether they are still valid in the given state of the system. This of course means sacrificing the advantages of the compilation approach in terms of local reasoning and reduced complexity. Fortunately, we can do better than checking all compiled axioms each time we perform reasoning.

The first possible improvement is to move away from an active checking for changes towards a mechanism where each local module remembers and records changes made to it. We can also think of a system where

498 individual modules actively notify other modules of changes to its local knowledge. This frees us from doing a
 499 complete check of the compiled knowledge and allows us to concentrate on these parts of the knowledge that
 500 actually were subject to changes.

501 The second improvement is in terms of an analysis of the impact a change in another module actually has
 502 on compiled knowledge. This is important as in real-world scenarios it turns out that a large part of the
 503 changes do not really affect the logical theory but are rather changes to the syntactic representation or changes
 504 in the naming of concepts and relations. While the latter have to be propagated to the definitions of the exter-
 505 nal concepts, they do not actually affect the compiled subsumption relations. Further, even if the logical theory
 506 underlying the ontology is affected by a change, this does not mean that it affects the compiled subsumption
 507 relations. This means that we have to find ways to distinguish changes that do have an impact on compiled
 508 relations from those that do not have an impact. In fact, the choice to restrict the language admissible in
 509 the definitions of external concepts allows us to precisely characterize these kinds of changes.

510 In the following, we concentrate on the analysis of the impact of changes on the validity of compiled
 511 subsumption relations. We first give a characterization of harmless and harmful changes. Here harmless
 512 changes are those that do not have an effect on compiled subsumption relations. Harmful changes are
 513 changes that do have a potential influence on compiled knowledge. We present mechanisms for classifying
 514 changes as harmless or harmful based on a syntactic analysis of changes made to an ontology. Finally, we
 515 present a simple mechanism that uses this information to decide whether the knowledge in a module has to
 516 be re-compiled.

517 4.1. Determining harmless changes

518 As compiled knowledge reflects subsumption relations between a query concept and a disjunction of other
 519 query concepts a harmless change is a set of modifications to an ontology that does not change these subsump-
 520 tion relations. Finding harmless changes is therefore a matter of deciding whether the modifications affect the
 521 subsumption relation between a query concept and a disjunction of other query concepts. It is quite obvious
 522 that a complete decision procedure for this problem has the same complexity as general subsumption reason-
 523 ing in the modular ontology and does therefore not improve the situation. For this reason, we propose a sound
 524 but incomplete method that abstracts from the detailed definition of concepts and uses the semantic relation
 525 between the old and the new version of a concept in the following way.

526 The method considers every concept and relation in an ontology module that has been subject to a change.
 527 Assuming that C represents the concept under consideration before and C' the concept after the change, there
 528 are four ways in which the old version C may semantically relate to the new version C' :

- 529 (1) the meaning of a concept is not changed: $C \equiv C'$ (e.g. because the change was in another part of the
 530 ontology, or because it was only syntactical);
 531 (2) the meaning of a concept is changed in such a way that the concept becomes more general: $C \sqsubseteq C'$;
 532 (3) the meaning of a concept is changed in such a way that the concept becomes more specific: $C' \sqsubseteq C$;
 533 (4) the meaning of a concept is changed in such a way that there is no subsumption relationship between C
 534 and C' .

535
 536 We can define the semantic relation between different versions of the same concept based on the set of pos-
 537 sible interpretations of the old and the new ontology in the following way.

538 **Definition 2** (*Semantics of Change*). Let C and C' be two version of the same concepts in ontologies O and O'
 539 respectively, then we say that C is more general than C' ($C \sqsupseteq C'$) if and only if for all possible interpretations \mathcal{I}
 540 and for all $x \in \Delta^{\mathcal{I}}$ we have $\mathcal{I}, O \models x \in C^{\mathcal{I}}$ implies $\mathcal{I}, O' \models x \in C'^{\mathcal{I}}$. Analogously, we say that C is more specific
 541 than C' ($C \sqsubseteq C'$) if an only if for all possible interpretations \mathcal{I} and for all $x \in \Delta^{\mathcal{I}}$ we have $\mathcal{I}, O' \models x \in C'^{\mathcal{I}}$
 542 implies $\mathcal{I}, O \models x \in C^{\mathcal{I}}$.

543 The same list holds for relations that are subject to change. The next question is how these different kinds of
 544 semantic relations between the old and the new version of a concept influences compiled knowledge. In order
 545 to understand this influence, we have to look at the influence of changes on the interpretation of query con-

546 cepts. We take advantage of the fact that there is a very tight relation between changes in concepts of the exter-
547 nal ontology and implied changes to the query concepts using these concepts:

548 **Lemma 1** (monotonicity of effect). *Let $C \equiv \mathcal{F}_j : Q$ an external concept expression. Let $c(Q)$ be the set of all*
549 *concept names and $r(Q)$ the set of all relation names occurring in Q , let further $C \in c(Q)$ and $R \in r(Q)$ then*
550 *changing C has the same impact on the interpretation of Q as it has on the interpretation of C , in particular, we*
551 *have $C \sqsubseteq C' \Rightarrow Q \sqsubseteq Q'$ and $C' \sqsubseteq C \Rightarrow Q' \sqsubseteq Q$ where Q' is the query as being interpreted after changing C .*
552 *Analogously, a change of R has the same effect on Q .*

553 **Proof Sketch** We prove lemma 1 by structural induction over expressions in the sublanguage defined in Sec-
554 tion 2.4.2. For the *induction basis* $Q = A$ the lemma trivially holds. In the induction step, we prove for every
555 operator that the whole expression becomes more general (specific) if either a concept or a relation occurring
556 in the expression becomes more general (specific). For conjunction and disjunction this directly follows from
557 their correspondence to set operations on the interpretation domain. It remains to be shown that lemma 1 also
558 holds for expressions of the form $(\geq nR.C)$ and $(\geq nR^-.C)$ (existential quantifiers are a special case). For
559 $(\geq nR.C)$ and $C \sqsubseteq C'$ this holds because all R-successors of C in $C^{\mathcal{I}}$ are also in $C'^{\mathcal{I}}$. Therefore we have
560 $(\geq nr.C)^{\mathcal{I}} \supseteq (\geq nr.C')^{\mathcal{I}}$. Further, there are no R-successors in $C'^{\mathcal{I}}$ that are not in $C^{\mathcal{I}}$. Therefore we have
561 $(\geq nr.C')^{\mathcal{I}} \subseteq (\geq nr.C)^{\mathcal{I}}$. The same argument holds for r^- . For $R \sqsubseteq R'$ the argument is similar. This time,
562 the subsumption follows from the fact that there are less members of C that are potentially in the relation
563 R with objects in Q' .

564 We can exploit this relation between the interpretation of external concepts and the concept names in their
565 definitions in order to identify the effect of changes in the external ontology on the subsumption relations
566 between different query concepts. First of all, the above result directly generalizes to multiple changes with
567 the same effect, i.e. a query Q becomes more general (specific) or stays the same if none of the elements in
568 $c(Q) \cup r(Q)$ become more specific (general). Further, the subsumption relation between an external concept
569 C and the disjunction of other external concepts does not change if all concepts in the disjunction become
570 more general or if the concept C becomes more specific. Combining these two observations, we derive the fol-
571 lowing characterization of harmless change.

572 **Theorem 2** (harmless change). *Let $C_0 \equiv \mathcal{F}_j : D_0, C_1 \equiv \mathcal{F}_j : D_1, \dots, C_m \equiv \mathcal{F}_j : D_m$ be external concept defini-*
573 *tions such that $\mathfrak{I} \models C_0 \sqsubseteq C_1 \sqcup \dots \sqcup C_m$, then a change is harmless with respect to the subsumption relation above*
574 *if:*

- 575 • $X' \sqsubseteq X$ for all $X \in c(D_0) \cup r(D_0)$,
576 • $X' \supseteq X$ for all $X \in c(D_i) \cup r(D_i), i = 1, \dots, m$

577 *Note again that the implication does not hold in the opposite direction.*

578 The theorem provides us with a correct but incomplete method for deciding whether a change is harmless given
579 that we know the semantic relation between the old and the new definition of concepts and relations that were
580 subject to changes. This method is a very basic version of the underlying idea of assessing the impact of changes.
581 We can think of more complete versions of the method that use a deeper analysis of the structure of the concept
582 expressions involved. Our experiences are, however, that this basic heuristic already covers most cases that occur
583 in practice, especially, because the definition above includes cases where most of the concepts stay unchanged.

584 4.2. Characterizing changes

585 Now that we are able to determine the consequence of changes in the concept hierarchy on the integrity of
586 the mapping, we still need to know what the effect of specific modifications on the interpretation of a concept
587 is (i.e. whether it becomes more general or more specific). As our goal is to determine the integrity of mappings
588 without having to do classification, we describe what theoretically could happen to a concept as result of a

Table 2

Some modifications to an ontology and their effects on the classification of concepts in the hierarchy

Operation	Effect
Add a role restriction to concept C	C : Specialized
<i>Complex</i> : change the superconcept of concept C to a concept lower in the hierarchy	C : Specialized
<i>Complex</i> : restrict the range of a role R (effect on all C that have a restriction on R)	R : Specialized, C : Specialized
Remove a superconcept relation of a concept C	C : Generalized
Change the concept definition of C from primitive to defined	C : Generalized
Add a concept definition A	C : Unknown
<i>Complex</i> : add a (not further specified) subconcept A of C	C : No effect
Define a role R as functional	R : Specialized

589 modification in the ontology. To do so, we have listed all possible change operations to an ontology according
 590 to the OWL² knowledge model in the same style as done in [35]. The list of operations is in principle extend-
 591 able to other knowledge models.

592 The list of change operations consists of two types of operations: (1) *atomic change operations*, such as *add*
 593 *range restriction* or *delete subconcept relation* and (2) *complex change operations*, which consist of multiple
 594 atomic operations and/or incorporate some additional knowledge. Complex changes are often more useful
 595 to specify effects than the atomic changes, as they incorporate some of the semantic consequences. For exam-
 596 ple, for operations like *concept moved up*, or *domain enlarged*, we can specify the effect more accurately than
 597 for the atomic operations *superconcept changed* and *domain modified*.³ Atomic changes can be detected at a
 598 structural level, i.e. by comparing the old and new definition of a concept, and are therefore computationally
 599 cheap with a linear complexity. To identify complex changes, we also need to take some of the semantic rela-
 600 tions in the ontology into account. This makes the complexity of the identification of complex changes poten-
 601 tially as bad as determining subsumption in *SHIQ*, i.e. Exp-Time Complete. However, in practice many
 602 complex changes can be detected at a structural level, e.g. by looking at explicitly stated subclass relations.

603 Table 2 contains some examples of operations and their effect on the classification of concepts. The table
 604 only shows a few examples, although our full ontology of change operations contains around 120 operations.
 605 This number is not fixed, as new complex changes can be defined. A snapshot of the change ontology can be
 606 found online.⁴ The specification of effects is not complete, in the sense that it describes “worst-case” scenarios,
 607 and that for some operations the effect is “unknown” (i.e. unpredictable). If the method’s output is
 608 “unknown” this means that in order to determine the semantic relation between the two versions we would
 609 have to perform logical reasoning. In contrast to [37] who provide complete semantics of changes, we prefer
 610 to use heuristics in order to avoid expensive reasoning about the impact of changes. By restricting the change
 611 detection to changes that can be detected at a structural level, the complexity of our change detection heuristic
 612 is linear.

613 4.3. Update management

614 With the elements that we described in this section, we now have a complete procedure to determine
 615 whether compiled knowledge in an ontology module is still valid when the external ontology modules are
 616 changed. The complete procedure is as follows. For each external concept C :

- 617 (1) determine the changes that are performed in the external ontology (e.g. by using the record of changes);
- 618 (2) heuristically determine the effect of the changes on the interpretation of the concepts and relations
 619 (where multiple changes to a concept or relation that have an opposite effect lead to the effect
 620 “unknown”);

² See <http://www.w3.org/TR/owl-features/>.

³ For a complete list, see [36].

⁴ <http://ontoview.org/changes/2/1>.

- 621 (3) create a list of all concept and relation names that occur in the external concept expression D for each
 622 concept C in the subsuming parts of the implied subsumption relations;
 623 (4) check whether all concepts and relations in this list remained unchanged or became more general;
 624 (5) create a list of all concept- and relation names that occur in the external concept expression D for each
 625 concept C in the subsumed parts of the implied subsumption relations;
 626 (6) check whether all concepts and relation in this list are unchanged or became more specific.
 627

628 In cases where we cannot guarantee that integrity is preserved, we recompute and re-compile the implied
 629 subsumption statements. We thus restore integrity and make correct local reasoning possible.

630 **Algorithm 2.** Update

```

631 Require: Ontology Module  $M$ 
632 Require: Ontology Module  $M_j$ 
633 forall compiled axioms  $C_1 \sqsubseteq D_1 \sqcup \dots \sqcup D_m$  in  $M^c$  do
634   forall  $X \in c(C) \cup r(C)$  do
635     if effect on  $X$  is 'generalized' or 'unknown' then
636        $M^c := \text{Recompile}(M, M_j)$ 
637     end if
638   end for
639 for all  $D_i, i = 1, \dots, m$  do
640   forall  $X \in c(D_i) \cup r(D_i)$  do
641     if effect on  $X$  is 'specialized' or 'unknown' then
642        $M^c := \text{Recompile}(M, M_j)$ 
643     end if
644   end for
645 end for
646 end for
647 end for
648 end for
649 end for
650 end for
651 end for
652 end for
653 end for
654
```

655 We describe the procedure in a more structured way in [Algorithm 2](#). The algorithm triggers a (re-)compilation step only if it is required in order to resume integrity. Otherwise no action is taken, because the previously compiled knowledge is still valid. In principle, all the steps can be automated. A tool that helps to automate steps (1) and (2) is described in [38]. This tool will compare two versions of an ontology and derive the list of change operations that is necessary to transform the one into the other.

660 The worst-time complexity of the update procedure once the effects of changes are known is linear in the number of concept and relation names occurring in compiled subsumption statements (in the worst case the effect of a change on every concept and relation name in the compiled axioms has to be checked). In practice, we expect the number of compiled axioms to be relatively small leading to a quite efficient procedure.

664 5. Application in a case study

665 In order to support the claims made about the advantage of modular ontologies, we applied our model in a small case study that has been carried out in the course of the WonderWeb project.⁵ Our main intention was to show that the update-management procedure presented in the last section can be used to avoid the computation of subsumption relations in many cases. For this purpose, we defined a small example ontology using mappings to an ontology in the human resource (HR) domain. We used the changes that occurred in the HR ontology during the different steps of the case study and determined the impact on our example ontology. Besides this, the case study provides us with examples of implied subsumption some of which are non-trivial but likely to occur in real-life situations.

⁵ See <http://wonderweb.semanticweb.org>.

673 5.1. The WonderWeb case study

674 WonderWeb is a EU/IST project that ran from 2002 till 2004. The aim of the project was to develop and
 675 demonstrate the infrastructure required for the large-scale deployment of ontologies as the foundation for the
 676 Semantic Web. In the context of this project the *Descriptive Ontology for Linguistic and Cognitive Engineering*
 677 (DOLCE) [39] has been developed. DOLCE is the first module in a library of foundational ontologies. The
 678 taxonomy of the most basic categories consists of concepts such as “location”, “social agent”, “event”, “par-

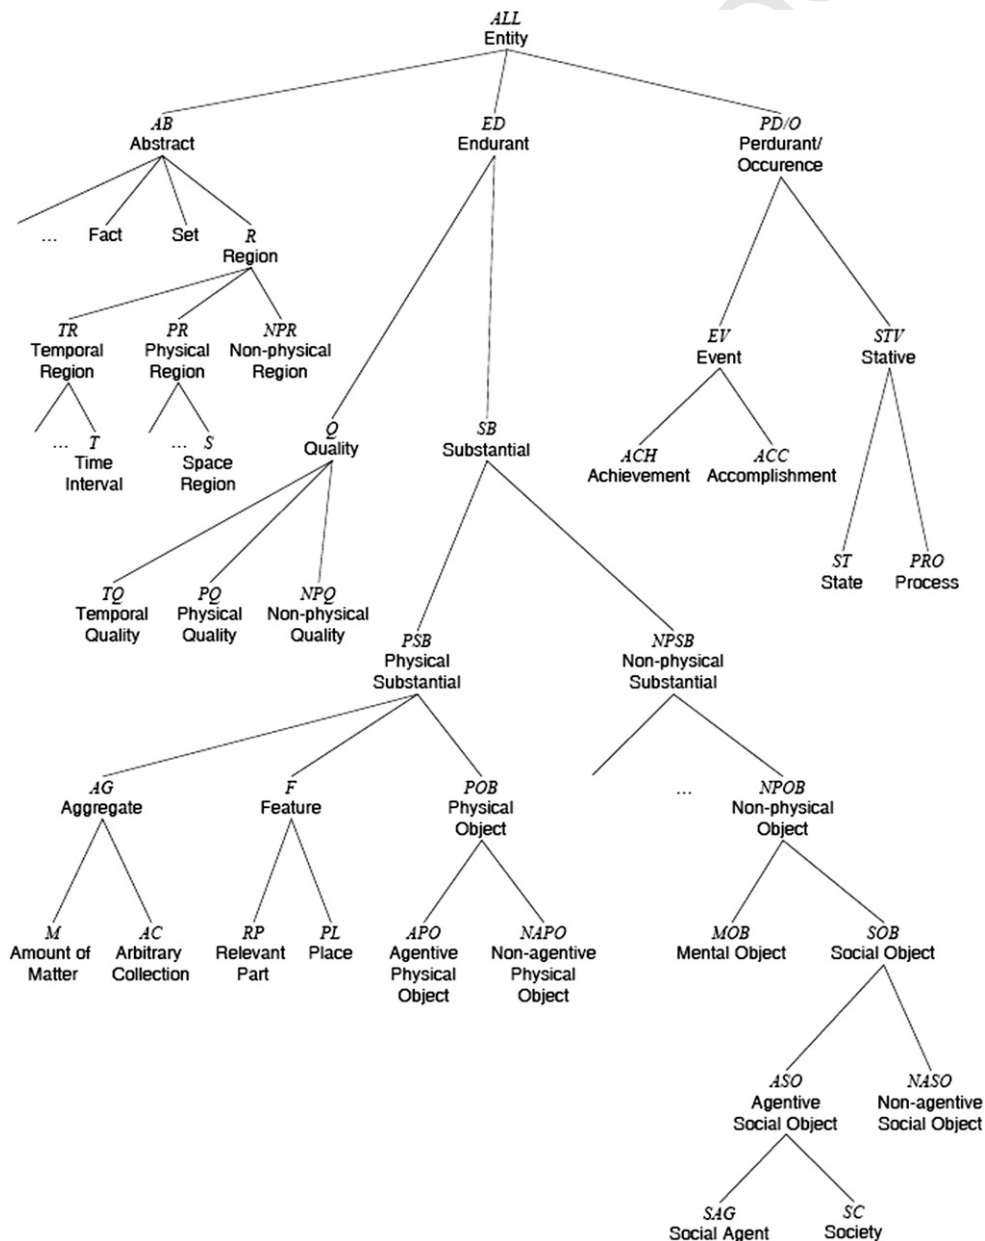


Fig. 1. Parts of the DOLCE upper level ontology (taken from [39]).

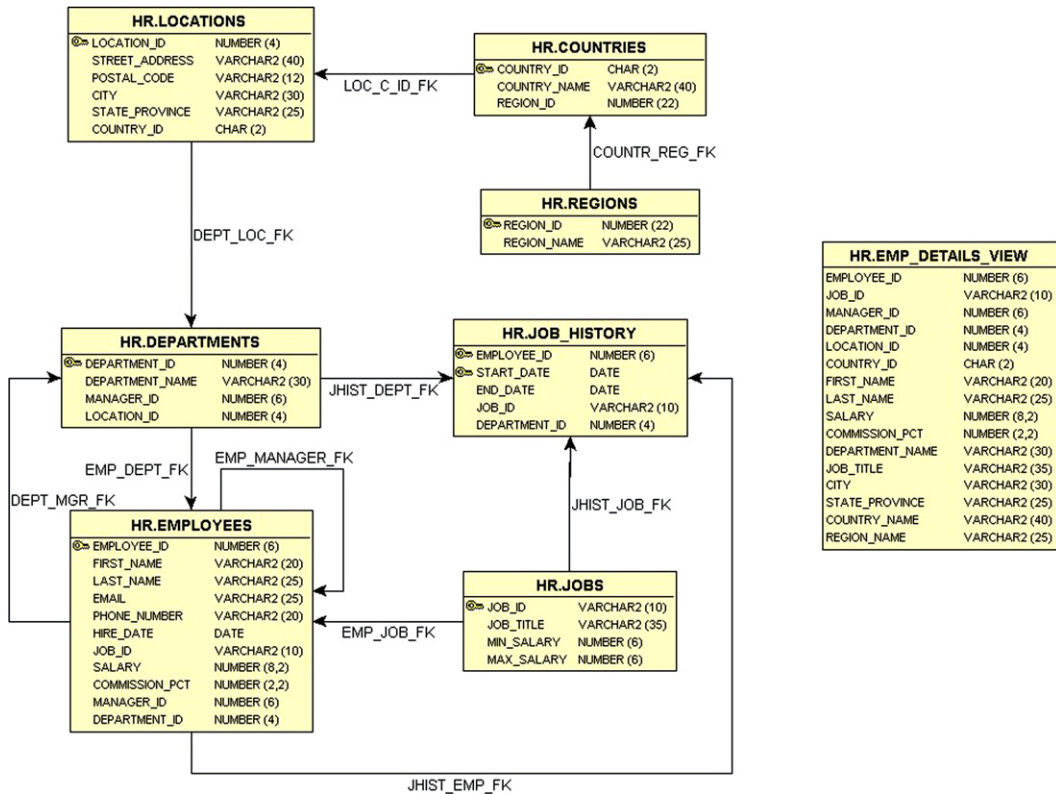


Fig. 2. Schema of the human resource database.

679 ticipant” and “region”. Fig. 1 shows the core part of the DOLCE ontology, the complete DOLCE ontology
 680 can be found online.⁶

681 One of the roles of foundational ontologies in general and DOLCE in particular is to clarify hidden
 682 assumptions underlying existing ontologies or linguistic resources by manually mapping existing categories
 683 into the categories defined in the foundational ontology [40]. Inconsistencies in the combined ontology point
 684 to modeling errors or wrong assumptions in the original ontology. Solving the inconsistencies will improve the
 685 quality of the initial ontology. The methodology around DOLCE describes a three step procedure for map-
 686 ping existing ontologies to DOLCE: alignment, refinement and tidying.

687 The case study that has been carried out in the project integrates different methods in the life-cycle of an
 688 ontology that is used on the Semantic Web, i.e. ontology creation, ontology refinement and ontology deploy-
 689 ment. The case study starts from an existing database schema in the human resource (HR) domain. The initial
 690 database schema is shown in Fig. 2. A first version of an ontology is created by a tool that automatically con-
 691 verts a database schema into an ontology [41]. In the next phase, the quality of the ontology is improved by
 692 manually mapping this ontology to DOLCE. First, the HR ontology is aligned with DOLCE, and in several
 693 successive steps the resulting ontology is further refined. During this process, the ontology changes continu-
 694 ously, which causes problems when other ontologies refer to definitions in the evolving ontology.

695 The original HR ontology combined with DOLCE is referred to as the DOLCE + HR ontology. In order
 696 to demonstrate the update management procedure, we created another ontology (which we call the *local ontol-*
 697 *ogy*) that uses terms and definitions from the evolving DOLCE + HR ontology (the *external ontology*). The

⁶ The main file can be found at <http://www.loa-cnr.it/ontologies/DOLCE-Lite.owl>, other files are referred from this file.

698 local ontology defines the concept *FulltimeEmployee* with a superconcept *Employee* and two subconcepts
699 *DepartmentMember* and *HeadOfDepartment*, using terms from the DOLCE + HR ontology.

700 The specific problem in our case is that the changes in the DOLCE + HR ontology could affect the reason-
701 ing in the local ontology. We want to be able to predict whether or not the reasoning in the local ontology is
702 still valid for specific changes in the external ontology.

703 The evolution of the DOLCE + HR ontology consisted of several steps. Each of these steps involves some
704 typical changes. We will briefly summarize them and show some changes that are typical for a specific step. In
705 Section 5.3, we discuss the effect of the changes described here on the integrity of implied subsumption rela-
706 tions in the local ontology.

- 707 • In the first step, the extracted HR ontology is aligned with the DOLCE foundational ontology, i.e. the con-
708 cepts and roles in the HR ontology are connected to concepts and roles in the DOLCE ontology via sub-
709 sumption relations. For example, the concept *Departments* from the HR ontology is made a subconcept of
710 *Social-Unit* in DOLCE.
- 711 • The refinement step involves a lot of changes. Some role restrictions are added, and some additional con-
712 cepts and roles are created to define the HR concepts more precisely. For example, the concept *Adminis-*
713 *trative-Unit* is introduced as a new subconcept of *Social-Unit*, and the concept *Departments* is made a
714 subconcept of it. Also, the range of the role *email* is restricted from *Abstract-Region* to its new subconcept
715 *Email* and the range of *manager-id* is specialized by restricting its range from *Employee* to the subconcept
716 *Manager*.
- 717 • In the next step, a number of concepts and roles are renamed to names that better reflect their meaning. For
718 example, *Departments* is renamed to *Department* (singular), and the two different variants of the role *man-*
719 *ager-id* are renamed to *employee-manager* and *department-manager*.
- 720 • In the final step, the tidying step, all roles and concepts that are not necessary any more are removed and
721 transformed into role restrictions. For example, the role *employee-email* is deleted from *Employee* and
722 replaced by an existential restriction in the concept *Person* on the role *abstract-location* to the concept *Email*.

723

724 5.2. Modularization in the case study

725 If we now consider the modularization in the case study, we have a local ontology with a concept hierarchy
726 that is built up by the following explicitly stated subsumption relations:

$$FulltimeEmployee \sqsubseteq Employee$$

$$DepartmentMember \sqsubseteq FulltimeEmployee$$

728 $HeadOfDepartment \sqsubseteq FulltimeEmployee$

729 This ontology introduces *FulltimeEmployee* as a new concept, not present in the case study ontology. Conse-
730 quently, this concept is only defined in terms of its relation to other concepts in the local ontology.

731 All other concepts are externally defined in terms of ontology based queries over the case study ontology.
732 The first external definition concerns the concept *Employee* that is equivalent to the *Employee* concept in the
733 case study ontology. This can be defined by the following trivial view:

735 $Employee \equiv HR : Employee$

736 Another concept that is externally defined is the *HeadOfDepartment* concept. We define it to be the set of all
737 instances that are in the range of the *department-manager* role. The definition of this view given below shows
738 that our approach is flexible enough to define concepts in terms of relations

740 $HeadOfDepartment \equiv HR : \exists department-manager. \top$

741 An example of a more complex external concept definition is the concept *DepartmentMember*, which is defined
742 using a query that consists of three conjuncts, claiming that a department member is an employee that is in the
743 *department-member* role with a department

$$\begin{aligned} \text{DepartmentMember} &\equiv \text{HR} : \text{Employee} \sqcap \\ &\exists \text{department}^- \text{member}^- . \text{Department} \end{aligned} \quad (4)$$

745

746 5.2.1. Implied subsumption relations

747 To allow for local reasoning, we need to determine the implied subsumption relations. If we now consider
748 logical reasoning about these external definitions, we immediately see that the definition of *Employee* sub-
749 sumes the definition of *DepartmentMember*, as the former occurs as part of the definition of the latter.

$$\models \text{DepartmentMember} \sqsubseteq \text{Employee} \quad (5)$$

753 At a first glance, there is no relation between the definition of *HeadOfDepartment* and the two other state-
754 ments as it does not use any of the concept or role names. However, when we use the background knowledge
755 provided by the case study ontology we can derive some implied subsumption relations. The reasoning is as
756 follows. Because the range of *department-manager* is set to *Employee* and the domain to *Department*, the def-
757 inition of *HeadOfDepartment* is equivalent to:

$$\text{Employee} \sqcap \exists \text{department}^- \text{manager}^- . \text{Department}$$

760 As we further know that *department-manager* is a subrole of *department-member*, we can derive the following
761 subsumption relation between the externally defined concepts:

$$\models \text{HeadOfDepartment} \sqsubseteq \text{Employee} \quad (6)$$

$$\models \text{HeadOfDepartment} \sqsubseteq \text{DepartmentMember} \quad (7)$$

765 Besides the subsumption relations between the external concepts themselves, we also need to determine the
766 subsumption relations between the external concepts and the disjunction of all combinations of other external
767 concepts. For our example, with $\text{HeadOfDepartment} \sqsubseteq \text{DepartmentMember} \sqsubseteq \text{Employee}$, this results in the fol-
768 lowing relations:

$$\models \text{Employee} \sqsupseteq \text{HeadOfDepartment} \sqcup \text{DepartmentMember} \quad (8)$$

$$\models \text{HeadOfDepartment} \sqsubseteq \text{Employee} \sqcup \text{DepartmentMember} \quad (9)$$

$$\models \text{DepartmentMember} \sqsubseteq \text{Employee} \sqcup \text{HeadOfDepartment} \quad (10)$$

772 When the relations (5)–(10) are added to the local ontology, it possible to do subsumption reasoning without
773 having to access the DOLCE + HR ontology any more.

774 5.3. Updating the models

775 We will now illustrate that the conclusions of the procedure are correct by studying the impact of some
776 changes mentioned in the bullet list in Section 5.1. We do not cover all changes, but discuss three typical cases.
777 For the other changes, a similar analysis could be done.

778 5.3.1. Example 1: The employee concept

779 For the *Employee* concept, the last step resulted in the removal of the *employee-email* relation. Our rules tell
780 us that this change makes the new version more general compared to its old version:

$$\text{Employee} \sqsubseteq \text{Employee}^{\text{e}}$$

783 According to our procedure, this should not be a problem because *employee* is in the “subsuming list”.

784 When we analyze this change, we see that it has an impact on the definition of the concept *Department-*
785 *Member* as it enlarges the set of objects allowed to take the first place in the has-member relation. This leads
786 to a new definition of *DepartmentMember* with $\text{DepartmentMember} \sqsubseteq \text{DepartmentMember}^{\text{e}}$. As *Department-*
787 *Member* was already more general than *HeadOfDepartment* and the *Employee* concept is not used in the def-
788 inition of the latter, the implied subsumption relation indeed still holds.

789 5.3.2. Example 2: The department-manager relation

790 In the second example, we have to deal with a change affecting a relation that is used in an external defini-
 791 tion. The relation *department-manager* is specialized by restricting its range to the concept *Manager* (which is
 792 a subconcept of *Employee*) making it a subrelation of its previous version (i.e. more specific):

794 $department_manager \sqsubseteq department_manager!$

795 According to our procedure, this change is harmful as *department-manager* is in the “subsuming list”.

796 Calculation and an analysis proves that this change indeed has impact on the definition of the concept
 797 *HeadOfDepartment*. Because the old version of the role *department-manager* is more general than the new
 798 one, it cannot be concluded that the old role (which is used in the external concept expression) is still a subrole
 799 of the *department-member* relation. Therefore, the implied subsumption relation *HeadOfDepartment*
 800 $\sqsubseteq DepartmentMember$ no longer holds. We have to recompute and compile the implied subsumption relations
 801 in order to guarantee integrity.

802 5.3.3. Example 3: The department concept

803 For the *Department* concept, different changes happened. For example, it has been made a subconcept of
 804 *Social-Unit*, it has been renamed from *Departments* and a relation has been removed. These changes have a
 805 contradictory effect, according to our heuristics, and as a consequence we cannot predict the relation between
 806 the old and the new versions. In this case our current heuristics leave us with no other option than recomput-
 807 ing the implied subsumption relations.

808 6. Summary

809 In this article, we discussed an infrastructure for the representation of and reasoning with modular ontol-
 810 ogies. The intention was to enhance the existing semantic web infrastructure with notions of modularization
 811 that have been proven useful in other areas of computer science, in particular in software engineering. We
 812 defined a set of requirements for modular ontologies that arise from expected benefits such as enhanced re-
 813 use and more efficient reasoning. Taking the requirements of loose coupling, self containment and integrity
 814 as a starting point, we defined a framework for modular ontologies providing the following contributions
 815 to the state of the art in ontology representation for the semantic web:

- 816 (1) We presented a formal model for describing dependencies between different ontologies. We proposed
 817 conjunctive queries for defining concepts using elements from another ontology and presented a
 818 model-based semantics in the spirit of Distributed Description Logics that provides us with a notion
 819 of logical consequence across different ontologies.
- 820 (2) We compared our model with the existing standard, i.e. the web ontology language OWL and showed
 821 that the OWL import facilities can easily be captured as a special case in our model. We further showed
 822 that our model provides additional expressiveness in particular with respect to modeling relations. In
 823 order to get a better idea of the improvements of our model over OWL, we investigated the formal prop-
 824 erties of inter module mappings, their impact on reasoning and their intuition.
- 825 (3) We described a method for detecting changes in an ontology and for assessing their impact. The main
 826 feature of this method is the derivation of conceptual changes from purely syntactic criteria. These con-
 827 ceptual changes in turn provide input for a semantical analysis of the effect on dependent ontologies, in
 828 particular on the validity of implied subsumption relations. We applied the method in a case study in the
 829 Wonder Web project and were able to determine the impact of changes without logical reasoning.
 830

831 7. Discussion

832 There are three major questions connected to the approach for reasoning and managing change in modular
 833 ontologies proposed in this article. The first is *feasibility* in terms of computational complexity. As mentioned
 834 above, we use a heuristic approach to tackle this problem, which raises the question about the *adequacy* of the

835 heuristics used. We argued that we chose a trade-off that works well in the context of OWL and typical seman-
836 tic web applications. This focus on a particular kind of representation finally raises the question of *generality*
837 of the approach. We discuss these three basic questions in the following.

838 7.1. Feasibility

839 Reasoning in modular ontologies is complex. We have shown that the complexity is essentially the same as
840 for reasoning in Classical Description Logics which are the basis for OWL. We cannot escape this complexity,
841 but we can move parts of the reasoning effort offline by compiling implied subsumption relations as described
842 in Section 3.2. This approach, however, is only feasible if there are phases where the offline computation nec-
843 essary to compile the implied relations can be done without affecting the performance of the system. Typically,
844 such computations are done ‘overnight’ when the system load can be assumed to be low. An alternative for
845 situations where this approach is not possible is to do the compilation on the fly. In particular, we can compile
846 implied subsumption relations whenever they are computed in order to answer a user query to the system. This
847 kind of ‘lazy compilation’ has the advantage that the enormous effort for compiling implied knowledge is done
848 as part of the normal reasoning process. In the beginning, users will not benefit much from this approach, but
849 the time savings increase with each query answered. In this way, we also prevent the compilation of knowledge
850 that is never used.

851 The main problem connected with the compilation approach, which is also a central aspect of this paper, is
852 the integrity of the compiled knowledge. In general compilation approaches only pay off if the computation
853 time saved by being able to use compiled knowledge is not larger than the effort of updating the compiled
854 knowledge. This means that compilation only makes sense in rather stable systems. In principle, we can
855 assume that knowledge on the terminological level as it is represented in ontologies is normally more stable
856 than instance data as normally found in databases. While changes to ontologies will occur less frequently they
857 can still have a significant impact on the system. For this reason, our work focussed on heuristics for efficiently
858 updating the system when changes occur. In this context, the feasibility of the approach relies on the adequacy
859 of the heuristics chosen.

860 7.2. Adequacy

861 Our method for detecting harmful changes is a conservative one. We basically identify changes that are
862 obviously not harmful and refrain from updating when only such changes occurs. The criteria used for this
863 purpose are rather weak as we do not consider the specific changes made to a concept but restrict our analysis
864 to the semantic relation between the old and the new version of the concept definition. The advantage of this
865 choice is the efficiency of the approach as the criterion for harmless changes can be checked in linear time with
866 respect to the number of concept names involved. We further weaken the method by not actually computing
867 the semantic relation between the old and the new version of the concept, but rather determine the relation
868 based on a set of change operations determined by syntactic analysis. This method also constitutes an incom-
869 plete heuristic as, for some changes, we cannot determine the relation between the old and the new version.
870 Again, the advantage of this choice lies in the efficiency of the approach, because it allows us to live completely
871 without Description Logic reasoning.

872 The question that arises is whether the degree of incompleteness of our approach is justified. Currently
873 there are no experimental results that support the usefulness of the heuristics, however, it is clear that the heu-
874 ristics cover a wide range of relevant cases. On the level of determining the semantic relation between the ver-
875 sions, the library of change operations covers all possible change operations and most of them have a known
876 effect that can be exploited in our approach. The cases in which the effect is not known can be assumed to be
877 the hard cases that will always require complete reasoning independent of the heuristics used. For this reason,
878 we believe that we cannot be much better on this level without falling back to classical subsumption reasoning.
879 On the level of determining harmless change, the heuristics used are quite weak as well. In particular we claim
880 that a change is only harmless if all concepts and relations satisfy certain properties. This can certainly be
881 relaxed if we invest more time in the analysis of the actual definitions of the concepts involved. For example,
882 changes in concepts that occur in both the subsumed and the subsuming concepts are not relevant due to the

883 monotonicity of the effect. We could also try to determine which parts of the definitions actually contribute to
884 the subsumption proof and consider changes in other parts of the definitions as harmless as well. These
885 improved heuristics still fit into the approach proposed and are subject of future work. For the time being,
886 we conclude that the general mechanisms are in place and that the heuristic described in this paper already
887 covers many relevant cases as we show in the examples from the WonderWeb case study. In principle, the
888 choice of the best heuristic will rely on the specific application scenario and in particular on the actual expres-
889 siveness used in the models.

890 7.3. Generality

891 A final point for discussion is the generality of the approach described. Throughout this paper, we based
892 our discussions on Description Logics as a representation language for ontologies, Distributed Description
893 Logics for providing the semantics of mappings as well as the equivalent of conjunctive queries for describing
894 relations between different modules. All of these choices are carefully made and are motivated by practical as
895 well as theoretical considerations. Probably the most uncontroversial choice is that of Description Logics for
896 encoding ontologies. In the context of semantic web research, Description Logics have become the primary
897 language for describing terminological knowledge mostly in terms of the Web Ontology Language OWL.
898 Our approach covers most of the expressiveness of OWL-DL with the exception of nominals. As a result, most
899 existing OWL ontologies will fit in our framework and could easily be turned into modular ontologies by add-
900 ing external concepts.

901 A choice that is less obvious is Distributed Description Logics as a basis for the semantics of mappings. In a
902 recent survey, we compared different approaches for describing mapping semantics [42]. One result of this
903 comparison was that Distributed Description Logics provide the highest degree of de-coupling between differ-
904 ent T-Boxes. This is important for our purposes as we want to support localized reasoning. A generalization of
905 Distributed Description Logics in terms of a distributed version of first order logic has been described by Ser-
906 afini and Ghididi [43]. We could have chosen this more general framework as the basis for our work, however,
907 the drawback of this is the lack of existing reasoning methods. Distributed Description Logics come with a
908 well investigated and implemented proof system that can be used to implement our approach.

909 The most controversial choice is to restrict the language that can be used to define external concepts. The
910 framework of Distributed Description Logic allows us to use arbitrary *SHIQ* expressions in the definitions.
911 A corresponding more general approach would have the same properties with respect to logical consequence,
912 compilation and local reasoning. The restriction to the equivalent of conjunctive queries was motivated by the
913 importance of the monotonicity property for the definition of update heuristics. This means that external con-
914 cepts can be defined using a more expressive language. This, however, would come at the price that implied
915 subsumption relations concerning this concept would have to be recomputed every time a change occurs. We
916 believe that the restriction proposed in this paper is reasonable as it allows the use of update heuristics and
917 also resembles view-based information integration, the dominant approach for describing mappings between
918 database schemata.

919 References

- 920 [1] S. Bechhofer, I. Horrocks, C. Goble, R. Stevens, OilEd: a reason-able ontology editor for the semantic web, in: F. Baader, G.
921 Brewka, T. Eiter (Eds.), KI 2001: Advances in Artificial Intelligence, Springer, 2001, pp. 396–408.
- 922 [2] N. Noy, R. Fergerson, M. Musen, The knowledge model of protege-2000: combining interoperability and flexibility (2000). URL
923 <citeseer.nj.nec.com/noy01knowledge.html>.
- 924 [3] I. Horrocks, The FaCT system, in: H. de Swart (Ed.), Automated Reasoning with Analytic Tableaux and Related Methods:
925 International Conference Tableaux'98, Lecture Notes in Artificial Intelligence, vol. 1397, Springer-Verlag, Berlin, 1998, pp. 307–312,
926 URL <download/1998/t98-paper.ps.gz>.
- 927 [4] V. Haarslev, R. Moller, Description of the RACER system and its applications, in: Proceedings of the Description Logics Workshop
928 DL-2001, Stanford, CA, 2001, pp. 132–142.
- 929 [5] J. Broekstra, A. Kampman, F. van Harmelen, Sesame: a generic architecture for storing and querying RDF and RDF schema, in: The
930 Semantic Web – ISWC 2002, Lecture Notes in Computer Science, vol. 2342, Springer, 2002, pp. 54–68.
- 931 [6] W. Farmer, J. Guttman, F. Thayer, Little theories, in: D. Kapur (Ed.), Proceedings of the Eleventh International Conference on
932 Automated Deduction, Lecture Notes in Computer Science, vol. 607, Springer-Verlag, 1992, pp. 567–581.

- 933 [7] P. Clark, J. Thompson, K. Barker, B. Porter, V. Chaudhri, A. Rodriguez, J. Thomere, S. Mishra, Y. Gil, P. Hayes, T. Reichherzer,
934 Knowledge entry as the graphical assembly of components, in: Proceedings of the 1st International Conference on Knowledge
935 Capture (K-Cap'01), 2001.
- 936 [8] E. Amir, S. McIlraith, Partition-based logical reasoning, in: 7th International Conference on Principles of Knowledge Representation
937 and Reasoning (KR'2000), 2000.
- 938 [9] S. McIlraith, E. Amir, Theorem proving with structured theories, in: B. Nebel (Ed.), Proceedings of IJCAI'01, Morgan Kaufmann,
939 San Mateo, 2001, pp. 624–634.
- 940 [10] S. Lauritzen, D. Spiegelhalter, Local computations with probabilities on graphical structures and their application to expert systems,
941 Journal of the Royal Statistical Society 50.
- 942 [11] A. Rector, Modularisation of domain ontologies implemented in description logics and related formalisms including OWL, in:
943 Proceedings of the 16th International FLAIRS Conference, AAAI, 2003.
- 944 [12] M. Buchheit, F.D.W. Nutt, A. Schaerf, Terminological systems revisited: Terminology = schema + views, in: Proceedings of the 12th
945 National Conference on Artificial Intelligence (AAAI-94), 1994.
- 946 [13] F. Giunchiglia, C. Ghidini, Local models semantics, or contextual reasoning = locality + compatibility, in: Proceedings of the Sixth
947 International Conference on Principles of Knowledge Representation and Reasoning (KR'98), Morgan Kaufmann, 1998, pp. 282–
948 289.
- 949 [14] A. Borgida, L. Serafini, Distributed description logics: directed domain correspondences in federated information sources, in: R.
950 Meersman, Z. Tari (Eds.), On the Move to Meaningful Internet Systems 2002: CoopIS, Doa and ODBase, LNCS, vol. 2519,
951 Springer-Verlag, 2002, pp. 36–53.
- 952 [15] D. Brickley, R. Guha, A. Layman, Resource description framework (RDF) schema specification, Working draft, W3C, August 1998.
953 <<http://www.w3c.org/TR/WD-rdf-schema>>.
- 954 [16] M. Dean, D. Connolly, F. van Harmelen, J. Hendler, I. Horrocks, D. McGuinness, P. Patel-Schneider, L. Stein, Web ontology
955 language (owl) reference version 1.0, Working draft, W3C, November 2002. <<http://www.w3.org/TR/owl-ref/>>.
- 956 [17] P. Bouquet, F. Giunchiglia, F. van Harmelen, L. Serafini, H. Stuckenschmidt, Contextualizing ontologies, Journal on Web Semantics
957 1 (4).
- 958 [18] R. Volz, A. Mdche, D. Oberle, Towards a modularized semantic web, in: Proceedings of the ECAI'02 Workshop on Ontologies and
959 Semantic Interoperability, 2002.
- 960 [19] R. Volz, D. Oberle, R. Studer, Views for light-weight web ontologies, in: Proceedings of the ACM Symposium on Applied Computing
961 SAC 2003, 2003.
- 962 [20] B. Cuenca-Grau, B. Parsia, E. Sirin, A. Kalyanpur, Modularity and web ontologies, in: P. Doherty, J. Mylopoulos, C. Welty (Eds.),
963 Proceedings of the International Conference on Knowledge Representation and Reasoning KR-06, 2006, pp. 198–209.
- 964 [21] J. Bao, D. Caragea, V. Honavar, On the semantics of linking and importing in modular ontologies, in: I. Cruz, S. Decker, D.
965 Allemang, C. Preist, D. Schwabe, P. Mika, M. Uschold, L. Aroyo (Eds.), Proceedings of the International Semantic Web Conference,
966 ISWC-06, Athens, GA, 2006, Lecture Notes in Computer Science, vol. 4273, pp. 72–86.
- 967 [22] A. Borgida, On the relative expressiveness of description logics and predicate logics, Artificial Intelligence 82 (1–2) (1996) 353–
968 367.
- 969 [23] D. Tsarkov, A. Riazanov, S. Bechhofer, I. Horrocks, Using vampire to reason with owl, in: Proceedings of the International Semantic
970 Web Conference, 2004, pp. 471–485.
- 971 [24] H. Stuckenschmidt, F. van Harmelen, Information sharing on the semantic Web, in: Advanced Information Processing, Springer-
972 Verlag, Berlin, Heidelberg.
- 973 [25] J. Gonzalez-Castillo, D. Trastour, C. Bartolini, Description logics for matchmaking of services, Technical Report HPL-2001-265, HP
974 Laboratories Bristol, 2001.
- 975 [26] L. Li, I. Horrocks, A software framework for matchmaking based on semantic web technology, International Journal of Electronic
976 Commerce 8 (4) (2004) 39–60.
- 977 [27] S. Bechhofer, I. Horrocks, D. Turi, The owl instance store: System description, in: Proceedings CADE-20, Lecture Notes in Computer
978 Science, Springer-Verlag, 2005.
- 979 [28] I. Horrocks, U. Sattler, S. Tobies, Reasoning with individuals for the description logic \mathcal{SHIQ} , in: Proceedings of the 17th
980 International Conference on Automated Deduction (CADE-17), Lecture Notes in Computer Science, Springer-Verlag, 2000.
- 981 [29] A. Borgida, L. Serafini, Distributed description logics: assimilating information from peer sources, Journal of Data Semantics 1
982 (2003) 153–184.
- 983 [30] L. Serafini, A. Borgida, A. Taminin, Aspects of distributed and modular ontology reasoning, in: Proceedings of the International Joint
984 Conference on Artificial Intelligence – IJCAI-05, Edinburgh, Scotland, 2005.
- 985 [31] F. Baader, D. Calvanese, D. McGuinness, D. Nardi, P. Patel-Schneider (Eds.), The Description Logic Handbook – Theory,
986 Implementation and Applications, Cambridge University Press, 2003.
- 987 [32] H. Stuckenschmidt, M. Klein, Integrity and change in modular ontologies, in: Proceedings of the 18th International Joint Conference
988 on Artificial Intelligence, Acapulco, Mexico, 2003. URL <<http://www.cs.vu.nl/mcaklein/papers/IJCAI03.pdf>>.
- 989 [33] S. Tobies, Complexity results and practical algorithms for logics in knowledge representation, PhD, thesis, LuFG Theoretical
990 Computer Science, RWTH-Aachen, 2001.
- 991 [34] L. Serafini, A. Taminin, DRAGO: distributed reasoning architecture for the semantic web, in: Proceedings of the Second European
992 Semantic Web Conference (ESWC'05), Springer-Verlag, 2005.
- 993 [35] J. Banerjee, W. Kim, H.-J. Kim, H.F. Korth, Semantics and implementation of schema evolution in object-oriented databases,
994 SIGMOD Record (Proc. Conf. on Management of Data) 16 (3) (1987) 311–322, URL <http://doi.acm.org/10.1145/38713.38748>.

- 24 *H. Stuckenschmidt, M. Klein / Data & Knowledge Engineering xxx (2007) xxx–xxx*
- 995 [36] M. Klein, Change management for distributed ontologies, Ph.D. thesis, Vrije Universiteit Amsterdam, August 2004. URL <<http://www.cs.vu.nl/mcaklein/thesis/>>.
- 996
- 997 [37] E. Franconi, F. Grandi, F. Mandreoli, A sematic approach to schema evolution and versioning in object-oriented databases, in: Proceeding of CL 2000, Lecture Notes in Artificial Intelligence, vol. 1861, Springer-Verlag, 2000, pp. 1048–1062.
- 998
- 999 [38] M. Klein, D. Fensel, A. Kiryakov, D. Ognyanov, Ontology versioning and change detection on the web, in: 13th International Conference on Knowledge Engineering and Knowledge Management (EKAW02), Sigüenza, Spain, 2002, LNCS, vol. 2473, p. 197ff. URL <<http://www.cs.vu.nl/mcaklein/papers/EKAW02.pdf>>.
- 1000
- 1001
- 1002 [39] A. Gangemi, N. Guarino, C. Masolo, A. Oltramari, L. Schneider, Sweetening ontologies with DOLCE, in: 13th International Conference on Knowledge Engineering and Knowledge Management (EKAW02), Sigüenza, Spain, 2002, Lecture Notes in Computer Science, vol. 2473, p. 166ff. URL <<http://link.springer.de/link/service/series/0558/bibs/2473/24730166.htm>>.
- 1003
- 1004
- 1005 [40] C. Masolo, S. Borgo, A. Gangemi, N. Guarino, A. Oltramari, Ontology library (final), Deliverable D18, EU/IST Project WonderWeb, December 2003. URL <<http://wonderweb.semanticweb.org/deliverables/documents/D18.pdf>>.
- 1006
- 1007 [41] R. Volz, D. Oberle, S. Staab, R. Studer, Ontolift prototype, Deliverable D11, EU/IST Project WonderWeb, 2002.
- 1008 [42] L. Serafini, H. Stuckenschmidt, H. Wache, A formal investigation of mapping languages for terminological knowledge, in: Proceedings of the 19th International Joint Conference on Artificial Intelligence – IJCAI05, Edingurgh, UK, 2005.
- 1009
- 1010 [43] C. Ghidini, L. Serafini, Distributed first order logic – revised semantics, Technical Report, ITC-irst, January 2005.
- 1011