Dynamic Change Evaluation for Ontology Evolution in the Semantic Web

I. Palmisano, V. Tamma
Dept. of Computer Science
University of Liverpool
L69 3BX, UK
{I.Palmisano, V.Tamma}@liverpool.ac.uk

L. Iannone
Dept. of Computer Science
University of Manchester
M13 9PL, UK
iannone@cs.man.ac.uk

T. R. Payne, P. Doran
Dept. of Computer Science
University of Liverpool
L69 3BX, UK
{T.R.Payne, P.Doran}@liverpool.ac.uk

Abstract

Changes in an ontology may have a disruptive impact on any system using it. This impact may depend on structural changes such as introduction or removal of concept definitions, or it may be related to a change in the expected performance of the reasoning tasks. As the number of systems using ontologies is expected to increase, and given the open nature of the Semantic Web, introduction of new ontologies and modifications to existing ones are to be expected. Dynamically handling such changes, without requiring human intervention, becomes crucial. This paper presents a framework that isolates groups of related axioms in an OWL ontology, so that a change in one or more axioms can be automatically localised to a part of the ontology.

1 Introduction

With the recent evolution and ubiquity of the World Wide Web (WWW), documents, systems and services are becoming increasingly decoupled, distributed, and decentralised [15]. To facilitate information exchange between diverse knowledge-based systems or agents[1], some form of consensus or agreement must exist over the choice of common vocabularies, or *ontologies* [17], that will be used. However, as such agents become situated in *open environments*, i.e. environments where distributed components may appear, change or disappear at any time, assumptions can no longer be made about content, interaction protocol, or even their availability or existence. Specifically, assumptions about a common domain ontology, adhered to by all agents, that can cater to the requirements of a diverse range of consumers and produces of services are no longer valid within this type of open environment. Thus, *ontology alignment*; i.e. the generation and use of mappings between two different ontologies describing the same domain, has provided a pragmatic (and principled) mechanism for facilitating interoperation between different ontologies, and consequently interaction between agents and services. Determining alignments, however, is a computationally expensive and semi-automated process that is traditionally performed offline. Whilst these alignments can be readily used when services are known a-priori, the use of alignments at runtime within an open environment requires mechanisms for locating existing alignments, and then selecting one (from a set of several possibilities) to facilitate service communication [9]. This process can be computationally heavy, and potentially redundant if a community of agents transact on a regular basis over a period of time. However, in these types of environments, no ontology can be expected to remain unchanged throughout its lifetime, but will *evolve*, and respond to changes in the environment. Such changes include: changes in the data represented in the ontologies; the need to accommodate the arrival of new agents; the improvement of the efficiency of repeated communication between a group of agents; or the dynamic determination of the alignment of their ontologies.

Traditional ontology evolution approaches [10] are semi-automated at best, and assume manual guidance from one (or several) domain experts. They are mainly static and rarely consider the use of estimates on the impact resulting from changes to an axiom on the whole ontology. Estimating this effect *a priori*, i.e. before performing the change itself, is even more crucial in open environments, where the agents' ability to acquire new capabilities and therefore to achieve new tasks (or answer new queries, in case of knowledge based systems) needs to be offset by the *cost* of the change in terms of employment of scarce resources [5], and with partial knowledge of the environment [12]. Whilst some studies have demonstrated that the addition of new axioms to the agent knowledge base increases its ability to achieve a task [14], few have attempted to estimate the impact of change, and none consider dynamic evolution [8].

In this paper we present an approach that evaluates the

[1] The term agent is here used loosely, to include intelligent agents, but also any computer system capable of autonomously performing a task.

impact of *change* on an ontology *a priori*, without using reasoning, but by heuristically determining the set of axioms in an ontology that will be affected by the change. This work assumes that the agent's decision making process is optimised to work with partial knowledge, and with limited computational resources. The approach determines the maximal scope of the effect of the modification performed in an ontology (*group*). The idea is that by identifying a group, agents can determine the impact of a proposed change, identify the reclassification costs involved, and therefore decide *a priori*, and without having to make use of reasoning, whether performing the change to their ontology is in their best interest. This work concentrates on introducing the formal definition of a group (Section 2.2), and the discussion of the properties that it exhibits (Section 3). Section 4 presents an empirical evaluation of the system, and the paper concludes by presenting final remarks.

2 Change Impact and Evaluation

OWL ontologies may undergo modifications for a variety of reasons; the simplest cases are those in which new individuals are added, or where the T-Box is updated to reflect changes in the domain or to increase the level of detail. Throughout this paper, we consider ontologies which are expressed in OWL DL [11], whose expressivity corresponds to the $\mathcal{SHOIN}(\mathcal{D})$ Description Logic (DL) [2].

Adding an individual to an A-Box is simpler than changing existing axioms or assertions, since all the previous inferences are still valid after the update, due to the monotonicity of DL, while the removal of axioms requires the removal of all inferences that the removed axioms entailed.

Another important source of modifications for an ontology comes from alignment with another ontology. Ontology alignment approaches provide a mechanism to facilitate interoperation between different ontologies, and thus interaction between agents and services.

Whilst the problem of ontology alignment has received much attention in recent years, with many methodologies and systems having been developed [4], little attention has been devoted to the way a knowledge base should cope with the additions introduced by alignments to another ontology. Often the solution consists of a new knowledge base, containing the two original ontologies and the alignment axioms that link them together. This approach has the advantage of being simple to implement, but in an open environment this may cause the knowledge base to grow uncontrollably, causing the performances to degrade beyond acceptable limits.

To illustrate how an alignment between two ontologies can affect a knowledge base, let us consider an example, that will be subsequently used throughout the paper:

Two agents, A_O and A_K, commit to two different ontologies O and K respectively (the T-Box of K is depicted in Figure 1), which differ on the definition of a specific concept[2], named C_2 in O, and which is not present in K. By means of an alignment technique, a concept C_1 in K is discovered as the best alignment for C_2, with the relation between them being an axiom $Ax = C_1 \subseteq C_2$. This enables agent A_O to issue queries against the knowledge base of agent A_K. For example, by asking for instances of C_2, the result will contain all the instances of C_1; this has the clear advantage of returning valid results to A_O, while issuing the same query without previous alignment would return no results.

2.1 Ontology Modification Impact

In the above example, K may be affected by the alignment, in that some axioms may be added. For example, let us consider the case in which the definition of C_2 (Def_{C_2}) is added to K, with Ax acting as a connection between Def_{C_2} and the original K, as depicted in Figure 1. The impact of adding this axiom to K has on the future performances of K cannot be easily predicted. The simplest heuristic is based on the assumption that $K \sqcup Ax \sqcup Def_{C_2}$ will, on average, behave in the same way as its expressivity class, i.e., the ontologies that use the same set of DL constructors to define their concepts. More informed heuristics are under development in literature [18].

If Def_{C_2} introduces new DL constructors, the expressivity will increase and the future performance for some reasoner over this ontology will decrease accordingly. Due to space limitations, it is not possible to illustrate in detail the different DL constructors and related complexity results here; however, more details can be found in [2]. Let us consider the modification consisting of the addition of Ax and Def_{C_2} to a knowledge base K, as depicted in Figure 1. K has expressivity $\mathcal{AL}$; by adding Ax and Def_{C_2} to K, the expressivity of this knowledge base will change to $\mathcal{ALC}$, raising then the computational complexity in the worst case from PTIME to PSPACE. In a larger knowledge base, there is the possibility that such an increase in expressivity may be restricted to a portion of the knowledge base, and in Section 4 we illustrate this for real world ontologies. The challenge is therefore to identify what portion of an ontology will be affected by the inclusion of an axiom. To address this challenge, we present an approach that considers each axiom and its interactions with the rest of the ontology.

The analysis that can be carried out is necessarily an estimate, since the aim is to evaluate the cost of a change *before* the change takes place, and a complete evaluation is only possible *after* the change has taken place. There-

[2]All concepts whose definition is not given are intended to be DL atomic concepts

fore, the impact of a change is measured using the very simple heuristic based on the number of axioms and assertions contained in the section of the knowledge base affected by the change, weighed with the expressivity of the knowledge base. Given the need for a computationally light way to assess the impact, the grouping framework does not try to find the exact set of axioms impacted by a change, but uses simple syntactic relations to compute candidate *groups*. Details on how the size of groups might be reduced, at the expense of the required effort, are given in Section 4.2.

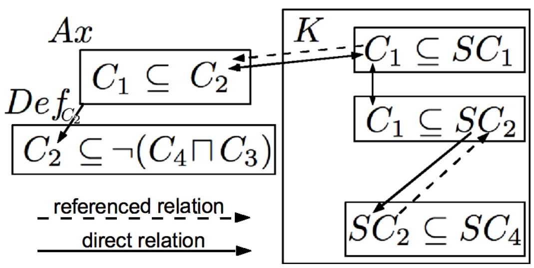

Figure 1. The group rooted at C_1 contains the T-Box K

2.2 Axiom and Group Definitions

DL ontologies typically contain a T-Box (DL axioms that define relations between concepts and roles) and an A-Box (DL axioms about individuals, also called assertions). In the grouping framework, an $axiom$ is the representation of a DL axiom or assertion; let us introduce an $axiom\ A$ for Ax (introduced in Section 2.1):

- A represents the abstraction over an OWL axiom or assertion, which has a *signature*, and is represented as a node in a directed multigraph whose edges represent *relations* between $axioms$;
- a *signature* s is the set of named and unnamed concepts and roles mentioned in an $axiom$; A has $s = \{C_1, C_2\}$;
- a *main concept* or *main role* m (also called *main node*)[3] is the concept or role being defined by the $axiom$[4]; m_A is C_1 in the example;
- *relations* between $axioms$ correspond to intersections between their *signatures* and *main nodes*; they can be:
  - **direct**: a *direct* relation between an $axiom\ A$ and an $axiom\ B$ exists if m_B belongs to the s_A; two $axioms$ with the same m have a bidirectional *direct* relation;
  - **indirect**: an *indirect* relation between two $axioms\ A$ and B holds when s_A overlaps with s_B (e.g. $D_1 \subseteq \exists\ R.C$ and $D_2 \subseteq \exists\ R.D$ share a reference to the role R); an *indirect relation* is bidirectional;
  - **referenced**: a *referenced* relation is the inverse of a *direct* relation; such a relation is implicitly defined also for *indirect* relations, in which case it is bidirectional as the *indirect relation*.

In DL, a signature is defined as the disjoint union **S** = **R** ⊎ **C** ⊎ **I** of *role names* (R), *concept names* (C) and *nominals* (I) appearing in an ontology; however, the definition given above to describe single axioms includes anonymous concepts, that in DL terms are only syntactic structures. This reflects more closely the abstract syntax for the OWL language[5] and simplifies the implementation, enabling a complete abstraction from the underlying RDF translation. Concept and roles defined in the language namespaces, i.e. OWL, RDF[6], RDFS[7] and XML Schema Datatype[8], are ignored when computing signatures.

If we consider O as the set of all the $axioms$ in an ontology $\mathcal{O}$, we can define three graphs based on the three kinds of relations, where: O_d is the graph $< O,\ DR >$ where DR is the set of edges that represent *direct* relations; O_i is the graph $< O,\ IR >$ where IR is the set of edges that represent *indirect* relations; O_r is the graph $< O,\ RR >$ where RR is the set of edges that represent the *referenced* relations. We define a group rooted at an axiom as follows:

Definition *A group G, rooted at an axiom A, is the set of $axioms$ resulting from the union of the sets of $axioms$ S_d, S_i and S_r explored during the exhaustive visit of O_d, O_i and O_r respectively, starting from A and following the relations.* □

Depending on the ontology, the size of a group can vary from a few axioms to the whole ontology; some results on real ontologies are reported in Section 4. To further illustrate this, consider the two simplest modifications to the knowledge base: the introduction of a new $axiom\ X$, and the removal of an $axiom\ Y$ already in the knowledge base.

In the first case, X may have relations to one or more $axioms$, and therefore to one or more groups:

[3] m may not be explicitly defined for all OWL axioms; for those cases where it is undefined, it is assumed to correspond to the common subject of all the triples involved in the $axiom$.

[4] An OWL axiom is represented internally as a set of RDF statements, where the standard RDF mapping for OWL is defined in *"OWL Web Ontology Language Semantics and Abstract Syntax Section 4: Mapping to RDF Graphs"* (`http://www.w3.org/TR/owl-semantics/mapping.html`).

[5] `http://www.w3.org/TR/owl-semantics/direct.html`

[6] `http://www.w3.org/RDF`

[7] `http://www.w3.org/TR/rdf-schema/`

[8] `http://www.w3.org/XML/Schema`

(a) X defines a new group, rooted at X, not influencing the existing groups. This is the case in which X has only *direct* outgoing relations to other axioms; the change thus only impacts the new group.

(b) X may be part of an existing group or act as a bridge between different groups. This happens when X has *direct* or *indirect* incoming relations from one or more group; the change impacts all involved groups.

In the second case, the removal of Y can have the following effects:

(a) Y may have relations to only one group and its removal does not break the group into subgroups; in this case, only the group rooted at Y is affected.

(b) Y may have been acting as a bridge between two or more groups; those groups will then be separate, and this will require the new groups to be reclassified.

Modifications to one or more axioms can be reduced to a composition of two kind of changes: the modification of a single $axiom$ is equivalent to the removal of the previous form of the $axiom$ and to the introduction of the new form, whilst the order in which multiple modifications are made is not relevant to the final result because of DL monotonicity.

Given a set of changes, it is therefore possible to select the affected groups and reduce the portion of the ontology that needs to be submitted to the reasoner. The advantage of using this framework is that determining the groups does not require a DL reasoner; it is in fact sufficient to take into account the syntactic characteristics of an ontology. In this way the heavy computational load of reasoning is not required until the knowledge base is actually used, thus enabling the use of grouping over large ontologies. On the other hand, the signature overlap between $axioms$ is not a guarantee that the $axioms$ are related, but it is only an indication; therefore, groups tend to be larger than strictly necessary, since not all relations mean that the related $axioms$ are used to entail relevant knowledge.

The grouping framework has been implemented in Java; Jena [3] and the SPARQL [9] implementation ARQ[10] have been used to perform the $axiom$ extraction. The OWL DL reasoning engine, Pellet [16], has been used to check the expressivity of a group, as reported in Section 4. In the following sections, we present theoretical considerations and empirical results.

3 Theoretical Properties of Groups

As defined in Section 2.2, a group consists of a set of related $axioms$. DL monotonicity guarantees that reasoning on a group G of a knowledge base K is *correct*, i.e., no $axiom$ that can be inferred from G cannot be inferred from K. This is referred to as *local correctness* in [6], where the concept of *uniform interpolant* [13, 19] is applied to prove the properties of a module extraction technique.

Local completeness would be a desirable property for a group, i.e., any DL axiom that can be entailed from K and which depends only on DL axioms in G should be entailed from G alone. However, to the best of our knowledge, this is still an open problem, since it depends on the ability to prove that G is a uniform interpolant for K. While a formal proof of this property for groups is not available, some relevant considerations are discussed below.

Let us consider the $\mathcal{ALC}$ DL and the corresponding tableaux expansion rules, as defined in [2]. Given an $axiom$ A, all *relevant* entailments that depend on A only depend on $axioms$ contained in G, the group rooted at A.

To support this claim, let us consider the expansion rules:

1. **$\sqcap$-rule**: if $\mathcal{A}$ contains $(C_1 \sqcap C_2)(x)$ and does not contain $C_1(x)$ and $C_2(x)$, then $\mathcal{A}_1 = \mathcal{A} \bigcup \{C_1(x), C_2(x)\}$.

2. **$\sqcup$-rule**: if $\mathcal{A}$ contains $(C_1 \sqcup C_2)(x)$ but neither $C_1(x)$ nor $C_2(x)$, then $\mathcal{A}_1 = \mathcal{A} \bigcup \{C_1(x)\}$, $\mathcal{A}_2 = \mathcal{A} \bigcup \{C_2(x)\}$.

3. **$\exists$-rule**: if $\mathcal{A}$ contains $(\exists R.C)(x)$, but there is no individual name z such that $C(z)$ and $R(x, z)$ are in $\mathcal{A}$, then $\mathcal{A}_1 = \mathcal{A} \bigcup \{C(y), R(x, y)\}$ where y is an individual name not occurring in $\mathcal{A}$.

4. **$\forall$-rule**: if $\mathcal{A}$ contains $(\forall R.C)(x)$ and $R(x, y)$, but it does not contain $C(y)$, then $\mathcal{A}_1 = \mathcal{A} \bigcup \{C(y)\}$.

If we consider rule (1), the possible situations are that:

(a) $\mathcal{A}$ contains the $axiom$ A=$X \equiv C_1 \sqcap C_2$, and the assertion $X(c)$ (also represented as an $axiom$ in the framework). These two are both contained in G, the group rooted at A, by means of a *direct* relation over X;

(b) $\mathcal{A}$ contains an assertion $C_1(y)$, which is relevant for the application of the $\sqcap$-rule. This assertion is linked to A through an *indirect* relation and therefore is included in G.

Any axiom or assertion relevant for rule (1) will be contained in G; therefore, restricting the input on which to apply the $\sqcap$-rule to G does not preclude any entailment in $\mathcal{A}$. The same approach works with rule (2); any assertion in $\mathcal{A}$ relevant for rule (2) will be included in G rooted at A=$X \equiv C_1 \sqcup C_2$.

For rules (3) and (4), the relation between the $axioms$ ($X \equiv \forall R.C$ and $X \equiv \exists R.C$) and the role assertions ($R(x, z)$ and $R(x, y)$) is an *indirect* relation based on the occurrence of R in both. The result in these cases is that

[9] SPARQL: Query language for RDF. W3C Recommendation, 15th January 2008. `http://www.w3.org/TR/rdf-sparql-query/`
[10] `http://jena.sourceforge.net/ARQ/`

all the relevant axioms and assertions are included in the groups G_1 rooted at A_1=$X \equiv \exists R.C$ and G_2 rooted at A_2=$X \equiv \forall R.C$ respectively. Therefore, it seems reasonable to assume that most of the relevant inferences that are entailed by $\mathcal{O}$ are also entailed by at least one group G computed over $\mathcal{O}$. Empirically, this has been proved true for the Galen fragment used in the evaluation. However, this result has still to be verified for full Galen, where the test is more difficult to perform due to the much larger size and complexity of the ontology. Extensions of these considerations that cover the higher expressivity of these ontologies (e.g. all of them include functional roles, which are not covered in the above) is ongoing at the time of writing.

3.1 Local Completeness

There are two main counterexamples for the *local completeness* of a group. The existence of an unsatisfiable axiom, such as $\top \subseteq \bot$, will extend beyond the local group and affect the whole ontology. Any ontology containing a contradiction may be used to entail any fact, and therefore its use in a real system is undesirable, and dubious at best. We therefore assume that rational agents will only use consistent ontologies to draw conclusions. The grouping framework is agnostic with respect to satisfiability in the input ontology, since no reasoning is used to compute the groups.

It is also possible to define a concept C that can be deduced as $C \equiv \top$; for example:

$$K = \{C \subseteq \top, C_1 \subseteq C, C_2 \subseteq C, C_1 \equiv \neg C_2\}$$

C therefore includes any concept defined in any group. A generic group G, however, would not include C and its subclasses unless some other axiom in K refers to C, C_1 or C_2, and therefore the axioms of the form $C_i \subseteq C$ for each C_i mentioned in K would not hold for the majority of groups. While this is a possible loss of information, on the other hand the construct that generates this effect is questionable, since in many cases it would be considered a modelling error in the input ontology; therefore, missing these deductions should not hamper the performances of most real world systems.

4 Practical Evaluation

A variety of real world ontologies were investigated to evaluate the effectiveness of axiom grouping and to isolate different groups. Of these, we report on OWL translation of the Galen ontology[11], and a fragment of Galen[12] which is smaller than the original and often used to test reasoners.

[11] http://www.co-ode.org/galen/

[12] http://www.daml.org/ontologies/400

Ontology	Expressivity	Groups	Size range
Galen full	$\mathcal{AL}$	72	1/45
	$\mathcal{ALC}$	14303	61723/63625
# axioms: 82030	$\mathcal{ALC}$	14303	61723/63625
# groups: 37408	$\mathcal{ALHIF}$	28	13/27
average size: 9649	$\mathcal{ALR}+\mathcal{HI}$	1	1
	$\mathcal{ALR}+\mathcal{HIF}$	313	53/1739
	$\mathcal{SHIF}$	24711	93/81815
Galen fragment	$\mathcal{AL}$	552	1/115
	$\mathcal{AL}(\mathcal{D})$	441	3/255
# axioms: 9915	$\mathcal{ALC}$	101	3/151
# groups: 4500	$\mathcal{ALC}(\mathcal{D})$	1257	4/7576
average size: 8529	$\mathcal{ALCH}$	1	14
	$\mathcal{ALCH}(\mathcal{D})$	29	20/7370
	$\mathcal{ALCHF}$	3	4/45
	$\mathcal{ALCHF}(\mathcal{D})$	95	16/7904
	$\mathcal{ALH}$	133	2/103
	$\mathcal{ALH}(\mathcal{D})$	7	4/79
	$\mathcal{ALHF}$	52	2/98
	$\mathcal{ALHF}(\mathcal{D})$	16	4/508
	$\mathcal{ALR}+\mathcal{H}$	7	2/29
	$\mathcal{ALR}+\mathcal{H}(\mathcal{D})$	1	25
	$\mathcal{ALR}+\mathcal{HF}$	29	5/170
	$\mathcal{ALR}+\mathcal{HF}(\mathcal{D})$	5	64/115
	$\mathcal{SH}(\mathcal{D})$	2	10/83
	$\mathcal{SHF}(\mathcal{D})$	427	47/9889

Table 1. Size and Expressivity Metrics

The measures considered include: the number of axioms in the ontology, the average number of axioms in a group and the expressivity of each group; for each expressivity level, the number of groups (duplicate groups are counted as one) and the size range are reported (Table 1).

4.1 Impact Estimation

Table 2 presents in more detail the results obtained on the Galen fragment ontology. Out of 9915 $axioms$, more than 4500 groups were computed, by selecting only $axioms$ with named concepts as m. The groups overlap very often, due to the high detail of the ontology, which is rich in definitions and has a deep role hierarchy. In fact, the number of distinct groups, i.e. groups not included in any other group, is only 7. One of these group is composed of a single axiom: $Ontology(galen)$; this represents the ontology itself, and does not interact with any axiom inside the ontology. The remaining 6 groups range in size from 23 axioms to 9889, and all of them but the smallest one have a common subset containing 7469 axioms; in the following, Int will be the intersection of the 5 largest groups, labelled G_2 to G_6, and will denote $G_i \setminus Int$ the set difference between one of these groups and Int. Only two groups are significantly larger than Int, G_5 and G_6, by 1923 and 2420 axioms respectively. In Table 2, the first table represents the size of the pairwise intersections between the groups, with the first column reporting the size (in number of axioms) of each of these groups; for readability, intersections of a set with itself are not reported, and neither are symmetric results.

The second table reports the size of Int; for each set, the size of the difference $G_i \setminus Int$ and its expressivity is given.

	# Axioms	G_3	G_4	G_5	G_6
G_1	23			6	23
G_2	7501	7485	7485	7497	7481
G_3	7488		7485	7487	7471
G_4	7503			7501	7485
G_5	9392				9376
G_6	9889				

	# Axioms	Expressivity
Int	7469	$\mathcal{ALC}$
$G_2 \setminus Int$	32	$\mathcal{AL}$
$G_3 \setminus Int$	19	$\mathcal{AL}$
$G_4 \setminus Int$	34	$\mathcal{AL}$
$G_5 \setminus Int$	1923	$\mathcal{SHF}$
$G_6 \setminus Int$	2420	$\mathcal{SHF}$

Table 2. Overlaps in Galen fragment grouping

The expressivity of the groups is $\mathcal{ALC}$ for the four smaller groups, and $\mathcal{SHF}$ for the two larger ones; the difference in expressivity is therefore related to roles ($\mathcal{SHF}$ is $\mathcal{ALC}$ plus transitive roles, functional roles and role hierarchy).

Given a modification to the ontology, now, it is possible to determine the part of the ontology that will be affected. Supposing the modification is the removal of an $axiom\ A$, the $axiom$ may be located at the intersection of the five largest groups, or it may be located within one specific group; using the following simple formula, computed considering the distinct groups containing A, a first approximation of the impact function may be computed:

$$I(A) = \sum_{A \in G} size(G) * expr(G)$$

where $size$ and $expr$ are defined as follows:

- $size()$: Given $(\mathcal{G})$ the set of all groups G over an ontology $\mathcal{O}$, $size(G) : \mathcal{G} \to \mathbb{N}$ is the number of $axioms$ contained in G.

- $expr()$: Given $(\mathcal{G})$ the set of all groups G over an ontology $\mathcal{O}$, $expr(G) : \mathcal{G} \to \mathbb{R}$ is the function that computes ec, the DL expressivity of G, and maps ec into a real numeric value.

The implementation of $expr()$ used to compute the results presented in this paper uses a list of four expressivity classes: EC = ($\mathcal{ALC}$, $\mathcal{SHIF}$, $\mathcal{SHOIN}$, $\mathcal{SROIQ}$) and a list of associated weights $ECWeights$=(0.25, 0.5, 0.75, 1); the DL expressivity ec of a group G is approximated to the least expressive $\mathcal{E}$ in EC that includes ec and the weight associated with $\mathcal{E}$ is returned.

The $expr$ function represents the weight that the reasoning has in evaluating the impact. Since worst case complexity can change with every allowed constructor in a DL, it is not simple to assign it a meaningful numeric value; as a first approximation, the values in $ECWeights$ have been chosen. Defining better approximations is one of the future developments of this work.

$A \in$:	Int	G_2	G_3	G_4	G_5	G_6
$I(A)$	47089	6086	6081	6088	14252	14580

Table 3. Results for impact computation

The values for $I(A)$ are reported in Table 3; they refer to all possible cases, i.e. A belonging to Int or to one of the other groups; on the basis of the comparison between the possible impacts, the decision making process may choose to accept the change involving A only in the cases in which the impact is smaller, which means the agent will accept the changes to $G_2 \setminus Int$, $G_3 \setminus Int$ and $G_4 \setminus Int$, therefore outside Int, while it will refuse changes to Int or to G_5 and G_6.

4.2 Related Work and Next Steps

Some recent efforts in literature have demonstrated that the addition of new axioms to an agent knowledge base increases its ability to achieve a task [14], but few efforts have attempted to estimate the impact of change, and not for dynamic evolution [8]. Sensoy and Yolum [14] propose a cooperation based approach in which agents exchange service descriptions, generated at runtime by one agent and passed to its "neighbours" so that, over time, the new service descriptions are common knowledge for the agents; however, agents are not required to incorporate the new service descriptions. The authors show, through simulations, that over time an agent tends to reuse service descriptions instead of creating new ones, i.e., including new knowledge helps the agent in reaching its goals. Sensoy and Yolum attack a slightly different topic than the one tackled by this paper: their work motivates the diffusion of knowledge between agents, proving that entities with heterogeneous knowledge will benefit from being able to exchange information and enrich their knowledge bases; however, no decision making criteria are defined. The work presented in this paper aims at devising an objective decision procedure for evaluating which modifications to a knowledge base will keep its growth in size and complexity in reasonable limits.

The current framework is susceptible to optimisations; as said in Section 2.2, grouping is based on the assumption that an overlap in the signature S means that the involved $axioms$ will entail useful knowledge when considered together, while this would not happen if they were to be placed in different groups; this assumption is not necessarily true, i.e. it is possible that some of the $axioms$ in a group do not participate in any inference, or the inferences they participate into can be drawn from a different set of $axioms$; this means that groups can be larger than necessary, thus reducing their utility in terms of size and expressivity reduction for impact evaluation. An alternative approach that does not cause the groups to grow more than necessary might be based on the work presented in [7];

Kalyanpur *et al* present an algorithm to find all explanations for a specific entailment; by applying this algorithm to entailments of a specific group, it is possible to verify which axioms are actually useful and which axioms can be left out of the group because they do not entail anything or the entailments in which they are involved can be inferred from a different set of axioms.

The work presented in [1] introduces the idea of partitioning axioms for First Order Logic, describing a partitioning algorithm and reasoning procedures based on *message-passing* with the individuated partitions, also providing correctness and completeness proof for the reasoning algorithms, based on Craig's *interpolation lemma*. As mentioned in Section 3, however, these results cannot be applied directly to DL formalisms.

5 Conclusions

This paper has introduced a framework for OWL knowledge base change evaluation through grouping, which can be useful in an open environment and can help an agent to make a rational choice when confronted with the possibility of changing its knowledge base. The paper presents a real world example, based on a fragment of the Galen ontology, to show the feasibility of rationally determining whether or not a specific change to the knowledge base may be accepted, using the proposed impact function.

An initial evaluation of the properties of a group has been presented, discussing the kind of entailed axioms that are not handled by the current definition of group. Ongoing developments are focused on the optimisation of grouping in order to reuse reasoning results between groups with overlapping sets of $axioms$, which actually cause a duplication of the reasoning effort. In addition, the theoretical properties of a group will be explored to provide more formal results characterising the behaviour of the framework, as well as an estimation of the feasibility of using more informed heuristics as evaluation measures.

Acknowledgements The research presented in this paper is supported by the EPSRC grant EP/D064287/1

References

[1] E. Amir and S. A. McIlraith. Partition-based logical reasoning. In *Principles of Knowledge Representation and Reasoning*, pages 389–400, 2000.

[2] F. Baader, D. Calvanese, D. McGuinness, D. Nardi, and P. Patel-Schneider, editors. *The Description Logic Handbook*. Cambridge University Press, 2003.

[3] J. J. Carroll, I. Dickinson, C. Dollin, D. Reynolds, A. Seaborne, and K. Wilkinson. Jena: implementing the Semantic Web recommendations. In S. I. Feldman, M. Uretsky, M. Najork, and C. E. Wills, editors, *Proc. of 13th Int. Conf. on World Wide Web - Alt. Track Papers & Posters, New York, NY, USA, May 17-20, pp. 74-83*. ACM, 2004.

[4] J. Euzenat and P. Shvaiko. *Ontology Matching*. Springer-Verlag, Heidelberg (DE), 2007.

[5] D. Fensel and F. van Harmelen. Unifying reasoning and search to web scale. *IEEE Internet Computing*, 11(2), 2007.

[6] B. C. Grau, B. Parsia, E. Sirin, and A. Kalyanpur. Modularity and web ontologies. In P. Doherty, J. Mylopoulos, and C. A. Welty, editors, *Proc. of the 10th Int. Conf. on Principles of Knowledge Representation and Reasoning, June 2-5*, pages 198–209. AAAI Press, 2006.

[7] A. Kalyanpur, B. Parsia, M. Horridge, and E. Sirin. Finding all justifications of owl dl entailments. In K. Aberer, K.-S. Choi, N. F. Noy, and (et al.), editors, *The Semantic Web, 6th Int. Semantic Web Conf., 2nd Asian Semantic Web Conf., ISWC + ASWC, Busan, Korea, November 11-15*, volume 4825 of *LNCS*, pages 267–280. Springer, 2007.

[8] A. Kalyanpur, B. Parsia, E. Sirin, and B. C. Grau. Repairing Unsatisfiable Concepts in OWL Ontologies. In Y. Sure and J. Domingue, editors, *3rd European Semantic Web Conf., Budva, Montenegro, June 11-14, 2006, Proc., LNCS 4011, pp. 170-184*. Springer.

[9] L. Laera, V. Tamma, J. Euzenat, T. Bench-Capon, and T. R. Payne. Reaching Agreement over Ontology Alignments. In *5th Int. Semantic Web Conf. (ISWC 2006)*, pages 371–384. Springer, November 2006.

[10] N. Noy and M. Klein. Ontology evolution: Not the same as schema evolution, 2003.

[11] P. F. Patel-Schneider, P. Hayes, and I. Horrocks. OWL Semantics and Abstract Syntax, 2004.

[12] T. R. Payne. Web services from an agent perspective. *IEEE Intelligent Systems*, 23(2):12–14, 2008.

[13] A. M. Pitts. On an interpretation of second order quantification in first order intuitionistic propositional logic. *J. Symb. Log.*, 57(1):33–52, 1992.

[14] M. Sensoy and P. Yolum. A cooperation-based approach for evolution of service ontologies. In Padgham, Parkes, Muller, and Parsons, editors, *Proc. of 7th Int. Conf. on Autonomous Agents and Multiagent Systems, May, 12-16, Estoril, Portugal*, pages 837–844, 2008.

[15] M. Singh and M. Huhns. *Service-Oriented Computing: Semantics, Processes, Agents*. Wiley, 2005.

[16] E. Sirin, B. Parsia, B. C. Grau, A. Kalyanpur, and Y. Katz. Pellet: A practical OWL-DL reasoner. *Journal of Web Semantics*, 5(2), 2007.

[17] R. Studer, V. R. Benjamins, and D. Fensel. Knowledge engineering: Principles and methods. *Data Knowledge Engineering*, 25(1-2):161–197, 1998.

[18] T. D. Wang and B. Parsia. Ontology Performance Profiling and Model Examination: First Steps. In K. A. et al., editor, *Proc. of 6th Int. Semantic Web Conf., ISWC 2007 + ASWC 2007, Busan, Korea, Nov. 11-15, 2007, LNCS 4825, pp. 595-608*. Springer.

[19] F. Wolter. Fusions of modal logics revisited. In M. Kracht, M. de Rijke, H. Wansing, and M. Zakharyaschev, editors, *Advances in Modal Logic*. CSLI, Stanford, CA, 1998.