

# A Novel Approach and Practical Algorithms for Ontology Integration

Giorgos Stoilos, David Geleta,  
Jetendr Shamdasani, and Mohammad Khodadadi

Babylon Health, London, SW3 3DD, UK  
`firstname.lastname@babylonhealth.com`

**Abstract.** Today a wealth of knowledge and data are distributed using Semantic Web standards. Especially in the (bio)medical domain several sources like SNOMED, NCI, FMA, and more are distributed in the form of OWL ontologies. These can be matched and integrated in order to create one large medical Knowledge Base. However, an important issue is that the structure of these ontologies may be profoundly different hence using the mappings as initially computed can lead to incoherences or changes in their original structure which may affect applications. In this paper we present a framework and novel approach for integrating independently developed ontologies. Starting from an initial seed ontology which may already be in use by an application, new sources are used to iteratively enrich and extend the seed one. To deal with structural incompatibilities we present a novel fine-grained approach which is based on mapping repair and alignment conservativity, formalise it and provide an exact as well as approximate but practical algorithms. Our framework has already been used to integrate a number of medical ontologies and support real-world healthcare services provided by Babylon Health. Finally, we also perform an experimental evaluation and compare with state-of-the-art ontology integration systems that take into account the structure and coherency of the integrated ontologies obtaining encouraging results.

**Keywords:** Ontology Integration, Ontology Matching, Ontology Alignment, Conservativity, Mapping Repair

## 1 Introduction

Today a wealth of knowledge and data are distributed using Semantic Web technologies and standards. For example, the Linked Open Vocabularies effort [22] contains more than 600 ontologies for various subjects like geography, multimedia, security, geometry, and more. Especially in the biomedical domain, a large number of ontologies have been developed during the previous decades like SNOMED,<sup>1</sup> NCI [5], UMLS,<sup>2</sup> the Disease ontology [16] and many more, while

<sup>1</sup> <https://www.snomed.org/>

<sup>2</sup> <https://uts.nlm.nih.gov/home.html>

BioPortal [15] is a repository of more than 600 biomedical ontologies. Identifying the common entities between these vocabularies and integrating them is beneficial for building ontology-based applications as one could unify complementary information that these vocabularies contain building a “complete” Knowledge Base (KB).

The problem of computing correspondences (mappings) between different ontologies is referred to as *ontology matching* or *alignment* [18]. A number of ontology matching systems have been developed in the past and have shown to behave well in practice. However, besides classes with their respective labels ontologies usually bring a class hierarchy and depending on how they have been conceptualised they may exhibit significant incompatibilities. For example, in NCI proteins are declared to be disjoint from anatomical structures whereas in FMA proteins are subclasses of anatomical structures. Hence, integrating them without taking into account their logical axioms may lead to many undesired consequences like unsatisfiable classes [11] and/or changes in their initial structure [6]. For these reasons the notions of *conservative alignment* [6, 19, 20] and *mapping repair* [12, 8] have been proposed in the literature. These notions dictate that the mappings should not alter the original ontology structure or introduce unsatisfiable concepts. If they do, then a so-called *violation* occurs which needs to be *repaired* by discarding some of the mappings.

Unfortunately, dropping mappings introduces another problem which is the increase of ambiguity and redundancy. For example, if one drops all mappings between NCI and FMA proteins (due to their structural incompatibilities), then the integrated ontology will contain at least two classes for the same real-world entity. Apart from an unnecessary increase in the size of the integrated ontology this introduces ambiguity and decreases interoperability between services that use classes from the KB. The problem becomes more acute if further sources are integrated in which case we may end up with multiple classes representing the same real-world entity.

An effort to construct a large medical KB by integrating existing medical sources recently started in Babylon Health.<sup>3</sup> The KB would serve as the backbone for healthcare services (diagnosis, drug prescription, and more) as well as for other tasks like medical text annotation, understanding, and reasoning. For these purposes a modular and highly configurable ontology integration framework was implemented which is using ontology mapping to discover correspondences between a new medical source and the current KB and enrich the latter with new medical knowledge. There were two major requirements in this effort. First, integrating new sources should not affect the behaviour of the services already functioning with the KB, hence its structure should not change when new sources are integrated. Moreover, the KB should not contain many entities with a large label overlap as this complicates text annotation tasks as well as doctors who are selecting classes from the KB for diagnostic purposes. To address these requirements our framework is using the notion of conservativity for tracking

---

<sup>3</sup> <https://www.babylonhealth.com/>

the structural changes, however, in order to repair them we propose a novel fine-grained approach which avoids dropping mappings as much as possible.

First, violations stemming from mappings of higher-multiplicity (i.e., those that map two entities from one ontology to the same entity in the other) are separated from the rest and both are treated differently since they are of different nature. The former are repaired by altering the mappings, however, the latter are repaired by dropping *axioms* from the new ontology. Our motivation is that services have already committed to the structure of the KB and parts of the new ontology that are in disagreement with this conceptualisation can be dropped. In addition, this approach helps reduce ambiguity and duplication of the integrated ontology as much as possible. Regarding violations on the structure of the new ontology, again a distinction between mappings of higher-multiplicity and the rest is made. The former are repaired by dropping mappings, however, the latter can be allowed since, as stated, the hierarchy of the new ontology is given low priority and various heuristics proposed in the literature can be used to guide this process. We formalised our framework using the notion of a (maximal) safe extension of a KB and provided an exact algorithm that is based on computing all repair plans [6].

Unfortunately, detecting all repair plans is known to be computationally very expensive [12, 8]. Consequently, we next present a concrete implementation of our framework which is using approximate but efficient algorithms for violation detection (all state-of-the-art systems are based on approximate algorithms). Our implementation has already been used to create a medical KB by integrating SNOMED, NCI, FMA, and CHV, and is currently under use within Babylon. We conclude the paper with an experimental evaluation and a comparison against state-of-the-art mapping repair systems obtaining encouraging results. In more detail, our implementation is currently the only one that can apply a general conservativity-based mapping repair strategy (not only mapping coherency detection) on such a large KB while the created KB contains far less distinct classes with overlapping labels (i.e., less ambiguity and duplication). In addition, it was the only approach for which no conservativity violations could be detected in the integrated KB.

## 2 Ontologies and Ontology Matching

For brevity reasons, throughout the paper we will use Description Logic notation. For a set of real numbers  $S$  we use  $\oplus S$  to denote the sum of its elements. For  $p$  an ontology prefix and  $C$  some class we sometimes write  $p:C$  to denote that  $C$  appears in ontology with prefix  $p$ . Hence, for IRI prefixes  $p_1 \neq p_2$ ,  $p_1:C$  and  $p_2:C$  denote different classes. For an ontology  $\mathcal{O}$  we use  $\text{Sig}(\mathcal{O})$  to denote the set of classes that appear in  $\mathcal{O}$ . Given an ontology  $\mathcal{O}$  we assume that all classes  $C$  in  $\mathcal{O}$  have at least one triple of the form  $\langle C \text{ skos:prefLabel } v \rangle$  and zero or more triples of the form  $\langle C \text{ skos:altLabel } v_i \rangle$ . For a given class  $C$  function  $\text{pref}(C)$  returns the string value  $v$  in the triple  $\langle C \text{ skos:prefLabel } v \rangle$ . An ontology is called *coherent* if every  $C \in \text{Sig}(\mathcal{O})$  with  $C \neq \perp$  is satisfiable.

In the literature, the notion of a Knowledge Base is almost identical to that of an ontology, i.e., a set of axioms describing the entities of a domain. In the following, we loosely use the term “Knowledge Base” ( $\mathcal{KB}$ ) to mean a possibly large ontology that has been created by integrating various other ontologies but, formally speaking, a  $\mathcal{KB}$  is an OWL ontology.

*Ontology matching* (or *ontology alignment*) is the process of discovering correspondences (mappings) between the entities of two ontologies  $\mathcal{O}_1$  and  $\mathcal{O}_2$ . To represent mappings we use the formulation presented in [12]. That is, a mapping between  $\mathcal{O}_1$  and  $\mathcal{O}_2$  is a 4-tuple of the form  $\langle C, D, \rho, n \rangle$ , where  $C \in \text{Sig}(\mathcal{O}_1)$ ,  $D \in \text{Sig}(\mathcal{O}_2)$ ,  $\rho \in \{\equiv, \supseteq, \sqsubseteq\}$  is the *mapping type*, and  $n \in (0, 1]$  is the confidence value of the mapping. Moreover, we interpret mappings as DL axioms—that is,  $\langle C, D, \rho, n \rangle$  can be seen as the axiom  $C \rho D$  with the degree attached as an annotation. Hence, for a mapping  $\langle C, D, \rho, n \rangle$  when we write  $\mathcal{O} \cup \{\langle C, D, \rho \rangle\}$  we mean  $\mathcal{O} \cup \{C \rho D\}$  while for a set of mappings  $\mathcal{M}$ ,  $\mathcal{O} \cup \mathcal{M}$  denotes the set  $\mathcal{O} \cup \{m \mid m \in \mathcal{M}\}$ . When not relevant and for simplicity we will often omit  $\rho$  and  $n$  and simply write  $\langle C, D \rangle$ . A *matcher* is an algorithm that takes as input two ontologies and returns a set of mappings.

### 3 An Ontology Integration Framework

Large KBs can be constructed by integrating existing, complementary, and possibly overlapping ontologies. For example, in the biomedical domain, ontologies for diseases, drugs, drug side-effects, genes, and so on, exist that can be integrated in order to build a large medical KB. However, before putting two sources together it would be beneficial to discover their overlapping parts and establish mappings between their equivalent entities.

*Example 1.* Consider an ontology-based medical application that is using the SNOMED ontology  $\mathcal{O}_{\text{snmd}}$  as a KB. Although SNOMED is a large and well-engineered ontology it is still missing medical information like textual definitions for all classes as well as relations between diseases and symptoms. For example, for class the notion of “Ewing Sarcoma” SNOMED only contains the axiom  $\text{snmd:EwingSarcoma} \sqsubseteq \text{snmd:Sarcoma}$  and no relations to signs or symptoms. In contrast, the NCI ontology  $\mathcal{O}_{\text{nci}}$  contains the following axiom about this disease:

$$\text{nci:EwingSarcoma} \sqsubseteq \exists \text{nci:mayHaveSymptom.nci:Fever}$$

We can use ontology matching to establish links between the related entities in  $\mathcal{O}_{\text{snmd}}$  and  $\mathcal{O}_{\text{nci}}$  and then integrate the two sources in order to enrich our KB. More precisely, using the labels of the aforementioned classes we can identify the following mappings:

$$\begin{aligned} m_1 &= \langle \text{snmd:EwingSarcoma}, \text{nci:EwingSarcoma}, \equiv \rangle \\ m_2 &= \langle \text{snmd:Fever}, \text{nci:Fever}, \equiv \rangle \end{aligned}$$

and hence replace our KB with  $\mathcal{O}'_{\text{snmd}} := \mathcal{O}_{\text{snmd}} \cup \mathcal{O}_{\text{nci}} \cup \{m_1, m_2\}$ . Then,  $\mathcal{O}'_{\text{snmd}}$  contains the knowledge that “Ewing sarcoma may have fever as a symptom”.  $\diamond$

Unfortunately, it is well-known that integrating ontologies using the initially computed mappings can lead to unexpected consequences like, introducing unsatisfiable classes [11] or structural changes to the input ontologies [6].

*Example 2.* Consider again the SNOMED and NCI ontologies. Both ontologies contain classes for the notion of “soft tissue disorder” and “epicondylitis”. Hence, it is reasonable for a matching algorithm to compute the following mappings:

$$\begin{aligned} m_1 &= \langle \text{snmd:SoftTissueDisorder}, \text{nci:SoftTissueDisorder}, \equiv \rangle \\ m_2 &= \langle \text{snmd:Epicondylitis}, \text{nci:Epicondylitis}, \equiv \rangle \end{aligned}$$

However, in NCI we have  $\mathcal{O}_{\text{nci}} \models \text{nci:Epicondylitis} \sqsubseteq \text{nci:SoftTissueDisorder}$  while in SNOMED  $\mathcal{O}_{\text{snmd}} \not\models \text{snmd:Epicondylitis} \sqsubseteq \text{snmd:SoftTissueDisorder}$ . Hence, in the integrated ontology we will have:

$$\mathcal{O}_{\text{snmd}} \cup \mathcal{O}_{\text{nci}} \cup \{m_1, m_2\} \models \text{snmd:Epicondylitis} \sqsubseteq \text{snmd:SoftTissueDisorder}$$

introducing a relation between classes of  $\mathcal{O}_{\text{snmd}}$  that did not originally hold and which can have a significant impact on the services of our application which are already based on the structure of  $\mathcal{O}_{\text{snmd}}$ .  $\diamond$

The amount of such structural changes can be captured by the notion of logical difference [9]. Like in [6] for performance reasons we also use an approximate version of logical difference formalised next.

**Definition 1 ([6]).** *Let  $A, B$  be atomic classes (including  $\top, \perp$ ), let  $\Sigma$  be a signature and let  $\mathcal{O}$  and  $\mathcal{O}'$  be two OWL 2 ontologies. The approximation of the  $\Sigma$ -deductive difference between  $\mathcal{O}$  and  $\mathcal{O}'$  (denoted  $\text{diff}_{\Sigma}^{\approx}(\mathcal{O}, \mathcal{O}')$ ) as the set of axioms of the form  $A \sqsubseteq B$  satisfying: (i)  $A, B \in \Sigma$ , (ii)  $\mathcal{O} \not\models A \sqsubseteq B$ , and (iii)  $\mathcal{O}' \models A \sqsubseteq B$ .*

Using logical difference, the notion of a *conservative alignment* has been proposed in the literature [6, 19, 20] which dictates that for two ontologies  $\mathcal{O}_1$  and  $\mathcal{O}_2$  and for  $\Sigma_1 = \text{Sig}(\mathcal{O}_1)$  and  $\Sigma_2 = \text{Sig}(\mathcal{O}_2)$  the set of mappings  $\mathcal{M}$  must be such that  $\text{diff}_{\Sigma_1}^{\approx}(\mathcal{O}_1, \mathcal{O}_1 \cup \mathcal{O}_2 \cup \mathcal{M})$  and  $\text{diff}_{\Sigma_2}^{\approx}(\mathcal{O}_2, \mathcal{O}_1 \cup \mathcal{O}_2 \cup \mathcal{M})$  are empty. An axiom belonging to either of these sets is called a (*conservativity*) *violation* and can be “repaired” by removing mappings from the initially computed sets.

Based on the above we have designed a Knowledge Base construction algorithm that is depicted in Algorithm 1. The algorithm accepts as input the current KB  $\mathcal{KB}$ , a new ontology  $\mathcal{O}$  which will be used to enrich  $\mathcal{KB}$  and a configuration `Config`. The configuration is used to tune and change various parameters like thresholds etc., many of which will be described in the rest of the paper. In brief, the algorithm first applies a set of matchers in order to compute a set of mappings between  $\mathcal{KB}$  and  $\mathcal{O}$  (lines 1–6). The set of matchers to be used is specified in the configuration object (`Config.Align.Matchers`) and each of them has a different weight assigned (`matcher.w`). After all matchers have finished, the mappings are aggregated and a threshold is applied (`Config.Align.thr`) in order to keep only mappings with a high confidence (lines 7–14). As mentioned previously, these

---

**Algorithm 1** KnowledgeBaseConstruction( $\mathcal{KB}, \mathcal{O}, \text{Config}$ )

---

**Input:** The current KB  $\mathcal{KB}$ , a new ontology  $\mathcal{O}$  and a configuration  $\text{Config}$ .

```
1: Mappings :=  $\emptyset$ 
2: for all matcher : Config.Align.Matchers do
3:   for all  $\langle C, D, \rho, n \rangle \in \text{matcher}(\mathcal{KB}, \mathcal{O})$  do
4:     Mappings := Mappings  $\cup \{ \langle C, D, \rho, n, \text{matcher} \rangle \}$ 
5:   end for
6: end for
7:  $\mathcal{M}_f := \emptyset$ 
8:  $w = \oplus \{ \text{matcher}.w \mid \text{matcher} \in \text{Config.Align.Matchers} \}$ 
9: for all  $\langle C, D, \rho, -, - \rangle \in \text{Mappings}$  such that no  $\langle C, D, \rho, n \rangle$  exists in  $\mathcal{M}_f$  do
10:   $n := \oplus \{ n_i \times \text{matcher}.w \mid \langle C, D, \rho, n_i, \text{matcher} \rangle \in \text{Mappings} \} / w$ 
11:  if  $n \geq \text{Config.Align.thr}$  then
12:     $\mathcal{M}_f := \mathcal{M}_f \cup \{ \langle C, D, \rho, n \rangle \}$ 
13:  end if
14: end for
15:  $\langle \mathcal{O}', \mathcal{M}_f \rangle := \text{postProcessNewOntoStructure}(\mathcal{KB}, \mathcal{O}, \mathcal{M}_f, \text{Config})$ 
16:  $\langle \mathcal{O}', \mathcal{M}_f \rangle := \text{postProcessKBStructure}(\mathcal{KB}, \mathcal{O}', \mathcal{M}_f, \text{Config})$ 
17: return  $\mathcal{KB} \cup \mathcal{O}' \cup \mathcal{M}_f$ 
```

---

mappings are still not the final ones since they may cause conservativity violations. These are handled by two functions, namely `postProcessNewOntoStructure` and `postProcessKBStructure` which are discussed in detail in the next section.

## 4 Safe Ontology Integration

The typical approach to resolve conservativity violations so far has been to remove mappings [6, 8, 19, 20]. However, this approach may introduce other issues like having distinct classes with a large overlap in their labels, hence introducing redundancy and ambiguity. Assume for instance, that in Example 2 we drop mapping  $m_2$ . Then, the integrated ontology will contain two different classes for the real-world notion of “epicondylitis” (i.e., `nci:Epicondylitis` and `snmd:Epicondylitis`) each with overlapping labels. Subsequently, a service that is using the former class internally cannot interoperate with a service that is using the latter as there is no axiom specifying that the two classes are actually the same.

Instead of removing mappings, another way to repair a violation is by removing axioms from one of the input ontologies.

*Example 3.* Consider again Example 2 where  $\mathcal{O}_{\text{snmd}}$  serves as the current version of the application KB. Instead of computing  $\mathcal{KB}_1^{\text{int}} := \mathcal{O}_{\text{snmd}} \cup \mathcal{O}_{\text{nci}} \cup \{m_1, m_2\}$  as in Example 2 assume that we compute the following:

$$\mathcal{KB}_2^{\text{int}} := \mathcal{KB}_1^{\text{int}} \setminus \{ \text{nci:Epicondylitis} \sqsubseteq \text{nci:SoftTissueDisorder} \}$$

Then, we have  $\mathcal{KB}_2^{\text{int}} \not\sqsubseteq \text{snmd:Epicondylitis} \sqsubseteq \text{snmd:SoftTissueDisorder}$  and hence  $\text{diff}_{\text{Sig}(\mathcal{O}_{\text{snmd}})}^{\approx}(\mathcal{O}_{\text{snmd}}, \mathcal{KB}_2^{\text{int}}) = \emptyset$  as desired.  $\diamond$

This approach is reasonable if we assume that an application is already using some Knowledge Base and the role of new ontologies is to enrich and extend it with new information but without altering its structure. Then, parts of the new ontology that cause violations can be dropped.

However, not all violations can be repaired by removing axioms from  $\mathcal{O}_2$ . This is the case for mappings of higher multiplicity, i.e., those that map two different classes of one ontology to the same class in the other.

*Example 4.* Consider again ontology  $\mathcal{O}_{\text{snmd}}$  and  $\mathcal{O}_{\text{nci}}$ . SNOMED contains classes *Eczema* and *AtopicDermatitis* whereas NCI contains class *Eczema* that also has “Atopic Dermatitis” as an alternative label. Hence, a matching algorithm could create two mappings of the form:

$$\begin{aligned} m_1 &= \langle \text{snmd:Eczema}, \text{nci:Eczema}, \equiv \rangle \\ m_2 &= \langle \text{snmd:AtopicDermatitis}, \text{nci:Eczema}, \equiv \rangle \end{aligned}$$

which imply that *snmd:Eczema* and *snmd:AtopicDermatitis* are equivalent although this is not the case in  $\mathcal{O}_{\text{snmd}}$ .  $\diamond$

In these cases it is clear that the only way to repair such violations is by altering the mapping set. One approach would be to drop one of the two mappings or perhaps even change their type from  $\equiv$  to  $\sqsubseteq$  or  $\supseteq$  and we argue that the choice is case dependent. In the previous example, we may decide that SNOMED is more granular than NCI in the sense that *Atopic Dermatitis* is a type of *Eczema* whereas the NCI term captures a more general notion. Hence, we may decide to change the mappings to  $\langle \text{snmd:Eczema}, \text{nci:Eczema}, \sqsubseteq \rangle$  and  $\langle \text{snmd:AtopicDermatitis}, \text{nci:Eczema}, \sqsubseteq \rangle$ . However, we may also conclude that the alternative labels in NCI don’t strictly denote synonym (alternative) terms for diseases but rather similar ones and hence decide to drop mapping  $m_2$ .

Based on the above we introduce the notion of a safe extension of an ontology.

**Definition 2.** *Let  $\mathcal{O}_1$  and  $\mathcal{O}_2$  be two ontologies and let  $\mathcal{M}$  be a set of mappings computed between them. The safe extension of  $\mathcal{O}_1$  w.r.t.  $\mathcal{O}_2, \mathcal{M}$  is a pair  $\langle \mathcal{O}', \mathcal{M}' \rangle$  such that  $\mathcal{O}' \subseteq \mathcal{O}_2, \mathcal{M}' \subseteq \mathcal{M}$  and  $\text{diff}_{\Sigma}^{\approx}(\mathcal{O}_1, \mathcal{O}_1 \cup \mathcal{O}' \cup \mathcal{M}') = \emptyset$  for  $\Sigma = \text{Sig}(\mathcal{O}_1)$ .*

The pair of an empty ontology and set of mappings  $(\langle \emptyset, \emptyset \rangle)$  is a trivial safe extension but one is usually interested in some maximal safe extension similar to the notion of diagnosis in mapping repair [8].

**Definition 3.** *Let  $\mathcal{O}_1$  and  $\mathcal{O}_2$  be two ontologies and let  $\mathcal{M}$  be a set of mappings computed between them. A safe extension  $\langle \mathcal{O}', \mathcal{M}' \rangle$  of  $\mathcal{O}_1$  w.r.t.  $\mathcal{O}_2, \mathcal{M}$  is called maximal if no safe extension  $\langle \mathcal{O}'', \mathcal{M}'' \rangle$  exists s.t. either  $\mathcal{O}'' \supset \mathcal{O}'$  or  $\mathcal{M}'' \supset \mathcal{M}'$ .*

Motivated by the above we have designed Algorithm 2 that accepts as input two ontologies  $\mathcal{O}_1, \mathcal{O}_2$  and returns a subset of  $\mathcal{O}_2$  and a subset of  $\mathcal{M}$  in an attempt to compute a safe extension. The algorithm first processes mappings of higher multiplicity w.r.t. entities in  $\mathcal{O}_1$  using function *disambiguate-m-1* whose properties are formalised next.

---

**Algorithm 2** `postProcessKBStructure( $\mathcal{O}_1, \mathcal{O}_2, \mathcal{M}, \text{Config}$ )`

---

**Input:** Two coherent ontologies  $\mathcal{O}_1, \mathcal{O}_2$ , a set of mappings  $\mathcal{M}$ , and a `Config` object.

- 1:  $\mathcal{M}_{m-1} := \{\langle C_i, D \rangle \mid \{\langle C_i, D \rangle, \langle C_j, D \rangle\} \subseteq \mathcal{M} \wedge C_i \neq C_j\}$ .
  - 2:  $\mathcal{M}' := \mathcal{M} \setminus \mathcal{M}_{m-1}$
  - 3: **for all**  $D \in \text{Sig}(\mathcal{O}_2)$  **do**
  - 4:      $\mathcal{M}' := \mathcal{M}' \cup \text{disambiguate-m-1}(\{\langle C_i, D \rangle \mid \langle C_i, D \rangle \in \mathcal{M}_{m-1}\}, \text{Config})$
  - 5: **end for**
  - 6:  $P := \text{allPlans}(\mathcal{O}_1 \cup \mathcal{O}_2 \cup \mathcal{M}', \mathcal{O}_2, \emptyset, \text{diff}_{\Sigma}^{\approx}(\mathcal{O}_1, \mathcal{O}_1 \cup \mathcal{O}_2 \cup \mathcal{M}'))$  where  $\Sigma = \text{Sig}(\mathcal{O}_1)$
  - 7: Pick some  $P \in P$  such that no  $P' \in P$  exists with  $P' \subset P$
  - 8: **return**  $\langle \mathcal{O}_2 \setminus P, \mathcal{M}' \rangle$
- 

**Definition 4.** Given a set of mappings  $\mathcal{M} = \{\langle C_1, D \rangle, \langle C_2, D \rangle, \dots, \langle C_n, D \rangle\}$  function `disambiguate-m-1` returns a set  $\mathcal{M}' \subseteq \mathcal{M}$  that satisfies the following property: it contains either a single mapping of the form  $\langle C_i, D, \equiv \rangle$  or only mappings of the form  $\langle C_i, D, \sqsupseteq \rangle$ .

Afterwards, the algorithm calls algorithm `allPlans` [6] passing the appropriate parameters in order to compute sets of axioms each of which has the following property: if it is removed from  $\mathcal{O}_2$  then all violations would be repaired. From all those sets the algorithm picks some minimal one (w.r.t.  $\subseteq$ ) and removes it from  $\mathcal{O}_2$ .

**Lemma 1.** Let  $\mathcal{O}_1$  and  $\mathcal{O}_2$  be two coherent ontologies and let  $\mathcal{M}$  be a set of mappings between them such that every class in  $\mathcal{O}_2$  is satisfiable in  $\mathcal{O}_1 \cup \mathcal{O}_2 \cup \mathcal{M}$ . When applied on  $\mathcal{O}_1, \mathcal{O}_2$  and  $\mathcal{M}$  Algorithm 2 returns a maximal safe extension of  $\mathcal{O}_1$  w.r.t.  $\mathcal{O}_2, \mathcal{M}$ .

Although, we are strict with respect to violations that are implied by the mappings to the structure of the KB, we can be more relaxed with respect to violations over the ontology that is being used for the enrichment. Several heuristics have been presented in the literature in order to decide which violations to allow and which to repair. A violation  $A \sqsubseteq B \in \text{diff}_{\text{Sig}(\mathcal{O}_2)}^{\approx}(\mathcal{O}_2, \mathcal{O}_1 \cup \mathcal{O}_2 \cup \mathcal{M})$  may be allowed if  $A$  and  $B$  are somehow semantically related, e.g., if  $A$  and  $B$  have a common descendant [19]. In contrast, a violation should be repaired if  $\mathcal{O}_2 \models A \sqsubseteq \neg B$ , i.e.,  $A$  and  $B$  are disjoint [11] or if the assumption of disjointness [13] can be applied to them—that is, if  $A$  and  $B$  are in different (distant) parts of the hierarchy of  $\mathcal{O}_2$  and hence we can assume that they are disjoint.

Motivated by the above we have designed Algorithm 3. Like before mappings of higher multiplicity are treated separately by function `disambiguate-1-m`. Afterwards, the algorithm iterates over all violations w.r.t. ontology  $\mathcal{O}_2$  and uses many of the aforementioned heuristics, like common descendants (line 8), unsatisfiability of classes (lines 9–12) and semantic or taxonomical similarity using function `semSim` and a pre-defined threshold `Config.Distance.thr` (lines 13–16) in order to decide to repair them or not. To figure out how to repair a violation algorithm `allPlans` is utilised again. This time  $\mathcal{M}'$  (and possibly  $\mathcal{O}_2$ ) instead of  $\mathcal{O}_1$  is passed as a second parameter since we don't want the plans to contain



---

**Algorithm 3** `postProcessNewOntoStructure`( $\mathcal{O}_1, \mathcal{O}_2, \mathcal{M}, \text{Config}$ )

---

**Input:** Ontologies  $\mathcal{O}_1, \mathcal{O}_2$ , set of mappings  $\mathcal{M}$ , and `Config` object.

```
1:  $\mathcal{M}_{1-m} := \{\langle C, D_i \rangle \mid \{\langle C, D_i \rangle, \langle C, D_j \rangle\} \subseteq \mathcal{M} \wedge D_i \neq D_j\}$ .
2:  $\mathcal{M}' := \mathcal{M} \setminus \mathcal{M}_{1-m}$ 
3: for all  $C \in \text{Sig}(\mathcal{O}_1)$  do
4:    $\mathcal{M}' := \mathcal{M}' \cup \text{disambiguate-1-m}(\{\langle C, D_i \rangle \mid \langle C, D_i \rangle \in \mathcal{M}_{1-m}\}, \text{Config})$ 
5: end for
6:  $\text{LDiff} := \text{diff}_{\Sigma}^{\approx}(\mathcal{O}_2, \mathcal{O}_1 \cup \mathcal{O}_2 \cup \mathcal{M}')$  where  $\Sigma = \text{Sig}(\mathcal{O}_2)$ 
7: for all  $A \sqsubseteq B \in \text{LDiff}$  do
8:   if no  $C$  such that  $\mathcal{O}_2 \models C \sqsubseteq A \sqcap B$  exists then
9:     if  $B = \perp$  then
10:       $\mathcal{O}_2^- := \{C \sqsubseteq \neg D \mid C \sqsubseteq \neg D \in \mathcal{O}_2\}$ 
11:       $P := \text{allPlans}(\mathcal{O}_1 \cup \mathcal{O}_2 \cup \mathcal{M}', \mathcal{O}_2^- \cup \mathcal{M}', \emptyset, \{A \sqsubseteq \perp\})$ 
12:       $\text{prune}(\mathcal{M}' \cup \mathcal{O}_2, P)$ 
13:     else if  $\text{semSim}(A, B) \leq \text{Config.Distance.thr}$  then
14:        $P := \text{allPlans}(\mathcal{O}_1 \cup \mathcal{O}_2 \cup \mathcal{M}', \mathcal{M}', \emptyset, \{A \sqsubseteq B\})$ 
15:        $\text{prune}(\mathcal{M}', P)$ 
16:     end if
17:   end if
18: end for
19: return  $\langle \mathcal{O}_2, \mathcal{M}' \rangle$ 
```

---

axioms from  $\mathcal{O}_1$ . Note that in case some class  $A$  is unsatisfiable in the integrated ontology either mappings from  $\mathcal{M}'$  or axioms from  $\mathcal{O}_2$  that lead to this unsatisfiability may be selected to be removed. This choice was motivated by the conceptual differences between NCI and FMA regarding anatomical structures and proteins which are disjoint in one ontology but semantically related in the other. The choice of which plan to pick to remove is again case dependent hence we abstract this away using the function `prune` which picks at-least one plan.

**Lemma 2.** *Let  $\mathcal{O}_1$  and  $\mathcal{O}_2$  be two coherent ontologies and let  $\mathcal{M}$  be a set of mappings between them. When applied on  $\mathcal{O}_1, \mathcal{O}_2$  and  $\mathcal{M}$  Algorithm 3 returns a pair  $\langle \mathcal{O}', \mathcal{M}' \rangle$  such that every class in  $\mathcal{O}'$  is satisfiable in  $\mathcal{O}_1 \cup \mathcal{O}' \cup \mathcal{M}'$ .*

Using Lemmas 1 and 2 we can show the main result of our paper.

**Theorem 1.** *Let  $\mathcal{KB}$  and  $\mathcal{O}$  be two coherent ontologies, let `Config` be some configuration and let  $\mathcal{KB}'$  be the output of Algorithm 1 when applied on  $\mathcal{KB}, \mathcal{O}$  and `Config`. Then the following hold:*

1.  $\text{diff}_{\Sigma}^{\approx}(\mathcal{KB}, \mathcal{KB}') = \emptyset$  where  $\Sigma = \text{Sig}(\mathcal{KB})$ .
2.  $\mathcal{KB}'$  is coherent.

## 5 Practical Algorithms

We have provided with concrete implementations of the algorithms and functions presented in the previous section. Regarding matching (lines 2-6 of Algorithm 1),

---

**Algorithm 4** `planApproximation( $\mathcal{O}_1, \mathcal{O}_2, \mathcal{M}$ )`

---

**Input:** Two ontologies  $\mathcal{O}_1$  and  $\mathcal{O}_2$  a set of mappings between them.

```
1: Exclusions :=  $\emptyset$ 
2: ConflictSets :=  $\{\{m_1, m_2\} \mid \mathcal{O}_1 \cup \mathcal{O}_2 \cup \{m_1, m_2\} \models_{\text{rdfs}} A \sqsubseteq B, \mathcal{O}_1 \not\models_{\text{rdfs}} A \sqsubseteq B\}$ 
3: for all  $\{\langle A, A' \rangle, \langle B, B' \rangle\} \in \text{ConflictSets}$  with  $\mathcal{O}_2 \models_{\text{rdfs}} A' \sqsubseteq B'$  do
4:   Exclusions := Exclusions  $\cup \{A' \sqsubseteq E \mid A' \sqsubseteq E \in \mathcal{O}_2, \mathcal{O}_2 \models_{\text{rdfs}} E \sqsubseteq B'\}$ 
5: end for
6: return Exclusions
```

---

two in-house label-based matchers have been implemented, namely `ExactLabelMatcher` and `FuzzyStringMatcher`. The former builds an inverted index of class labels [3] after some string normalisations, like removing possessive cases (e.g., Alzheimer’s) and singularisation [10], and matches ontologies using these indexes. The latter is based on the `ISub` string similarity metric [21]. Since this algorithm does not scale well on large inputs [3] it is mostly used for disambiguating higher-multiplicity mappings or if we wish to re-score subsets of mappings with low confidence degrees. In addition to these matchers, the state-of-the-art systems `AML` [4] and `LogMap` [7] can also be used in Algorithm 1.


Regarding functions `disambiguate-m-1` and `disambiguate-1-m` the following strategy has been implemented so far:

For a set of mappings  $\{\langle C_1, D \rangle, \langle C_2, D \rangle, \dots, \langle C_n, D \rangle\}$  and some real-value threshold `Config.Disamb.th`, if  $i \in [1, n]$  exists such that the following two conditions hold:

1. `ISub(pref( $C_i$ ), pref( $D$ )) > ISub(pref( $C_j$ ), pref( $D$ ))` for every  $j \neq i$  and
2. `ISub(pref( $C_i$ ), pref( $D$ ))  $\geq$  Config.Disamb.th`

then return  $\langle C_i, D \rangle$ . Similarly in function `disambiguate-1-m` for sets of mappings of the form  $\{\langle C, D_1 \rangle, \langle C, D_2 \rangle, \dots, \langle C, D_n \rangle\}$ .

A major practical consideration of Algorithms 2 and 3 is the call to algorithm `allPlans`. This algorithm does not scale in practice since it iterates over the power-set of the second parameter (i.e.,  $\mathcal{O}_2$ ) [6]. Consequently, as usually done in literature [8, 19] we are using approximations of plan computation and violation repair. More precisely, lines 6 and 7 in Algorithm 2 are replaced by the call  $P := \text{planApproximation}(\mathcal{O}_1, \mathcal{O}_2, \mathcal{M}')$  where the implementation of this function is given in Algorithm 4 and is inspired by the Alcomo repair algorithm [12, 8]. This algorithm is based on the assumption that logical differences of the form  $A \sqsubseteq B$  stem from *exactly* two mappings which map classes  $A$  and  $B$  for which  $\mathcal{O}_1 \not\models A \sqsubseteq B$  to classes  $A'$  and  $B'$  in  $\mathcal{O}_2$  for which a path of `SubClassOf` axioms in  $\mathcal{O}_2$  exists ( $\models_{\text{rdfs}}$ ), hence implying changes in the structure of  $\mathcal{O}_1$ . Although in theory this may not always be the case, most violations in practice do follow this pattern. For every such pair of mappings the algorithm picks to remove from  $\mathcal{O}_2$  some axiom of the form  $A' \sqsubseteq E$ , i.e., it tries in some sense to remove the “weakest” axiom from  $\mathcal{O}_2$ . This choice is motivated by belief revision and the *principle of minimal change* [1].



**Fig. 1.** Ontologies of Example 5 and resulting KBs; with bold we denote classes from  $\mathcal{O}_2$ .

*Example 5.* Consider for example the following two ontologies:

$$\begin{aligned}\mathcal{O}_1 &= \{D \sqsubseteq C, C \sqsubseteq B\} \\ \mathcal{O}_2 &= \{W \sqsubseteq Z, Z \sqsubseteq Y, Y \sqsubseteq X\}\end{aligned}$$

and assume the set of mappings  $\mathcal{M} = \{m_1, m_2\}$  where  $m_1 = \langle D, Y \rangle$  and  $m_2 = \langle B, W \rangle$ . Clearly, for  $\Sigma = \text{Sig}(\mathcal{O}_1)$  we have  $B \sqsubseteq D \in \text{diff}_{\Sigma}^{\approx}(\mathcal{O}_1, \mathcal{O}_1 \cup \mathcal{O}_2 \cup \mathcal{M})$  and we can repair this violation by either removing  $ax_1 = W \sqsubseteq Z$  or  $ax_2 = Z \sqsubseteq Y$ .

Ontologies  $\mathcal{O}_1$  and  $\mathcal{O}_2$  as well as KBs  $\mathcal{KB}_{ax_1}^{\text{int}} = \mathcal{O}_1 \cup \mathcal{O}_2 \setminus \{ax_1\} \cup \mathcal{M}$  and  $\mathcal{KB}_{ax_2}^{\text{int}} = \mathcal{O}_1 \cup \mathcal{O}_2 \setminus \{ax_2\} \cup \mathcal{M}$  are depicted graphically in Figure 1, where solid lines denote subclass relations, and dashed lines the two mappings. As we can see, although both integrated ontologies do not exhibit violations over  $\mathcal{O}_1$ , the two cases differ in the amount of changes they impose on the classes of  $\mathcal{O}_1$ . More precisely, for  $\mathbf{S}(\mathcal{O}_1, \mathcal{O}_2, \mathcal{M}) = \{A \sqsubseteq B \mid A \in \text{Sig}(\mathcal{O}_1), \mathcal{O}_1 \cup \mathcal{O}_2 \cup \mathcal{M} \models A \sqsubseteq B\}$  we have  $\mathbf{S}(\mathcal{O}_1, \mathcal{O}_2 \setminus \{ax_2\}, \mathcal{M}) \setminus \mathbf{S}(\mathcal{O}_1, \mathcal{O}_2 \setminus \{ax_1\}, \mathcal{M}) = \{B \sqsubseteq Z, C \sqsubseteq Z, D \sqsubseteq Z, D \sqsubseteq X\}$ . Indeed, in this scenario Algorithm 5 will compute  $\text{Exclusions} := \{ax_1\}$ .  $\diamond$

Following a similar approach, the block of lines 6–18 in Algorithm 3 is replaced by the steps depicted in Algorithm 5. Again these steps assume that violations stem from pairs of “conflicting” mappings like those mentioned above. Similarly to Algorithm 3, we are again using the heuristics of common descendants, disjoint classes and class similarity as a guide for repairing the violations; all entailments are checked using RDFS-entailment.

Using Algorithm 1 and the techniques presented above we have started building a large medical KB to be used within Babylon Health. We used the SNOMED January 2018 release (which contains 340K classes and 511K SubClassOf axioms) as a starting seed KB ( $\mathcal{KB}_1$ ) and have so far iteratively integrated the following ontologies: NCI version 17.12d (which contains 130K classes and 143K SubClassOf axioms), CHV latest version from 2011 (which contains 57K classes and 0 SubClassOf axioms) and FMA version 4.6.0 (which contains 104K classes and 255K SubClassOf axioms); all ontologies are from the official websites. As a matching algorithm we have so far used our ExactLabelMatcher. Statistics

---

**Algorithm 5** newOntologyApproximate( $\mathcal{O}_1, \mathcal{O}_2, \mathcal{M}$ )

---

```
1: ConflictSets :=  $\{\{m_1, m_2\} \mid \mathcal{O}_1 \cup \mathcal{O}_2 \cup \{m_1, m_2\} \models_{\text{rdfs}} A \sqsubseteq B, \mathcal{O}_2 \not\models A \sqsubseteq B\}$ 
2: for all  $\{\langle D_1, D'_1 \rangle, \langle D_2, D'_2 \rangle\} \in \text{ConflictSets}$  do
3:   if no  $D$  such that  $\mathcal{O}_2 \models_{\text{rdfs}} D \sqsubseteq D'_1 \sqcap D'_2$  exists then
4:     if  $D'_1 \sqsubseteq \neg D'_2 \in \mathcal{O}_2$  and  $C$  exist s.t.  $\mathcal{O}_1 \cup \mathcal{O}_2 \cup \mathcal{M}' \models_{\text{rdfs}} C \sqsubseteq D'_1 \sqcap D'_2$  then
5:       prune( $\mathcal{M}' \cup \mathcal{O}_2, \{\{\langle D_1, D'_1 \rangle, \langle D_2, D'_2 \rangle\}, \{D'_1 \sqsubseteq \neg D'_2\}\}$ )
6:     else if  $\text{semSim}(D'_1, D'_2) \leq \text{Config.Distance.thr}$  then
7:       prune( $\mathcal{M}', \{\{\langle D_1, D'_1 \rangle, \langle D_2, D'_2 \rangle\}\}$ )
8:     end if
9:   end if
10: end for
```

---

**Table 1.** Statistics about the KB after each integration/enrichment iteration.

	SNOMED	+NCI	+CHV	+FMA
Classes	340 995	429 241	429 241	524 837
Properties	93	124	124	219
SubClassOf Axioms	511 656	617 542	617 542	713 313
ObjPropAssertions	526 146	664 742	664 742	962 190
DataPropAssertions	543 416	946 801	1 043 874	1 211 459

about the KBs that we created after each integration are depicted in Table 1. CHV is a flat list of layman terms of medical concepts. From that ontology we only integrated label information for the classes in CHV that mapped to some class in the Babylon KB; hence only data type properties increased in the KB in that step. The final KB is currently under use by various services within Babylon and we are in the process of also integrating the following resources: SNOMED drug extensions, ICD-10, Read Codes, MeSH, and more.

## 6 Evaluation

We have conducted an experimental evaluation in order to assess the effectiveness of our approach for integrating ontologies and remedying conservativity violations. Using SNOMED as our initial Knowledge Base we once integrated NCI and then FMA (starting again from scratch). We used our ExactLabelMatcher once with and once without the last post-processing steps in Algorithm 1 (lines 15 and 16). In the following we call the former setting **bOWLing** and the latter **bOWLing<sub>n</sub>**. We used the latter setting as a “naive” baseline approach.

In addition, we also run Algorithm 1 using AML and two versions of LogMap called LogMap<sub>o</sub> and LogMap<sub>c</sub> in the following. AML and LogMap<sub>o</sub> repair mappings with respect to coherency, i.e., they only check for conservativity violations that lead to unsatisfiable classes. NCI contains 196 while FMA 33.5K disjoint classes axioms so this mapping repair is relevant. In contrast, LogMap<sub>c</sub> also checks for more general conservativity violations using the techniques presented in [19]. For all these systems we disabled the post-processing steps of Algorithm 1 in order to assess their mapping repair functionality. On the mapping

**Table 2.** Evaluation results

SNOMED+NCI						
	$ \mathcal{M} $	$ \mathcal{KB}^{\text{int}} $	$ \text{LDiff} $	Time	Loops	ambiguity
bOWLing <sub>n</sub>	30 675	677 939	t.o.	12.7	127	16 708
bOWLing <sub>n</sub> <sup>Alc</sup>	26 825	666 834	0.9m	35.9	100	17 177
bOWLing	19 258	638 702	0	12.2	0	7 810
LogMap <sub>o</sub>	27 967	664 837	1.7m	120.9	74	17 632
LogMap <sub>o</sub> <sup>Alc</sup>	27 763	664 354	1.5m	141.7	71	16 986
LogMap <sub>c</sub>	21 838	433 711	897	54.4	0	8,266
AML	32 623	635 876	t.o.	75.0	298	14 353
SNOMED+FMA						
	$ \mathcal{M} $	$ \mathcal{KB}^{\text{int}} $	$ \text{LDiff} $	Time	Loops	ambiguity
bOWLing <sub>n</sub>	8 809	614 728	240k	7.0	3	1 946
bOWLing <sub>n</sub> <sup>Alc</sup>	7 886	615 291	93k	76.2	1	2 000
bOWLing	8 176	608 060	0	27.9	0	1 440
LogMap <sub>o</sub>	7 334	615 252	117k	360.4	1	2 264
LogMap <sub>o</sub> <sup>Alc</sup>	6 986	615 689	57k	428.4	1	2 253
LogMap <sub>c</sub>	6 036	420 424	517	14 004.8	0	1 553

sets computed by bOWLing<sub>n</sub> and LogMap<sub>o</sub> we have also run Alcom [12] as a post-processing step. Alcom is not a general matcher but a mapping repair system that can be used as a post-processing step. In the following we denote these settings as bOWLing<sub>n</sub><sup>Alc</sup> and LogMap<sub>o</sub><sup>Alc</sup>.

Algorithm 1 did not terminate with AML and LogMap<sub>c</sub> after running for more than 16 hours. As a second attempt we fragmented the ontologies into modules (using a  $\sqsubseteq$ -reachability based algorithm starting from top-level classes) and fed these one by one to Algorithm 1. For NCI we extracted 53 while for FMA 6 modules. Even in this case AML did not terminate when integrating FMA.

Our results are summarised in Table 2 where we give the number of computed mappings ( $|\mathcal{M}|$ ), the number of SubClassOf axioms in the integrated ontology ( $|\mathcal{KB}^{\text{int}}|$ ), the number of axioms in  $\text{diff}_{\text{Sig}(\mathcal{KB})}^{\approx}(\mathcal{KB}, \mathcal{KB}^{\text{int}})$  (denoted by  $|\text{LDiff}|$  and with “m” denoting millions), and the time to compute  $\mathcal{KB}^{\text{int}}$  (in minutes). Due to the very large size of the KB LDiff cannot be computed by any OWL reasoner so we computed the RDFS-level differences by simply traversing the SubClassOf hierarchy of the KB. In addition, we have also computed the following:

- number of cycles of the form  $\{A_1 \sqsubseteq A_2, \dots, A_n \sqsubseteq A_1\} \subseteq \mathcal{KB}^{\text{int}}$ . From a semantic point of view such cycles are not problematic, however, they do complicate graph-based algorithms like hierarchy traversal, extracting paths and depth counting, hence it is a design decision in Babylon to avoid them; input ontologies contain no cycles.
- a notion of “ambiguity” which we defined as the number of times a label appears in two different classes of a given ontology. We also calculated this metric over the original SNOMED, NCI, and FMA ontologies in order to measure their level of ambiguity. We obtained 1 055, 4 873, and 282, respectively, e.g., in SNOMED 1 055 labels appear in multiple classes.

First thing to note from the table is that all systems compute mapping sets of comparable size with the exception of **bOWLing** on SNOMED+NCI which computes a smaller mapping sets. This is mostly due to functions **disambiguate-m-1** and **disambiguate-1-m** which prune mappings of higher-multiplicity. However, we should note that all mappings computed by this approach are one-to-one mappings, while in all other approaches from the roughly 27k mappings about 17k are actually one-to-one (i.e., fewer than those of **bOWLing**). The application of Alcom on the mapping sets does remove some mappings in an attempt to repair the sets while **LogMap<sub>c</sub>** that uses a general conservativity-based repairing approach also computes fewer mappings than **LogMap<sub>o</sub>**.

As expected, the ontology produced by **bOWLing** contains fewer axioms due to the axiom exclusion strategy implemented in line 16 of Algorithm 1 which drops about 30% of NCI axioms and 10% of FMA axioms. However, the gains from this approach are apparent when considering other computed metrics. More precisely, the integrated ontology produced by **bOWLing** contains no axioms in **LDiff** in contrast to even more than 1 million new ancestor classes in some of the other approaches. Moreover, there are no cycles and, finally, a very low degree of ambiguity taking also into account the initial ambiguity of these ontologies (see above). The use of Alcom as a post-processing step on **bOWLing<sub>n</sub>** and **LogMap<sub>o</sub>** does improve the numbers on these metrics, however, as it only focuses on coherency and not general conservativity it does not eliminate them completely. The only comparable approach is **LogMap<sub>c</sub>** which computes a KB without cycles. However, **LDiff** is still not empty and the approach of dropping mappings increases the ambiguity metric. Recall that we were only able to run **LogMap<sub>c</sub>** on the modules. Had it run on the whole ontology we believe the reported numbers would be higher since as one can note the integrated ontology in this module approach is also much smaller (almost 1/3 smaller). Finally, compared to all other systems our approach is much more scalable requiring a few minutes whereas in all other settings Algorithm 1 could take from one even up to 4 hours (even when restricted to the modules). Note that in some cases we could not compute **LDiff** even after 12 hours (t.o.).

## 7 Related Work and Conclusions

Constructing large Knowledge Bases (also called Knowledge Graphs these days) is a topic of intensive research and engineering the last years. The works [2, 17] focus on extracting medical facts from text and use ontologies like UMLS and SNOMED mostly as flat vocabularies for performing named entity disambiguation, text annotation and information extraction. Hence, the focus is not on merging the medical knowledge and “fusing” the hierarchies. Malacards [14] is an effort for constructing a large disease KB by integrating information from many existing disease ontologies. To identify the overlaps between different sources a label-based unification algorithm is used which is very similar to our Exact-LabelMatcher (labels are normalised, stemmed, singularised, etc. and a hash is created). However, this approach completely discards the hierarchy and the

axioms of the original ontologies and the final output is a flat non-ontological structure.

In the current paper we have studied the problem of building large KBs from existing ontologies by integrating them and retaining as much of their initial structure and axioms as possible. Starting with an initial ontology as a seed KB we use new ontologies to extend and enrich it in an iterative way. Overlaps are discovered using ontology matching algorithms and mappings are post-processed in order to preserve properties of the structures of the KB and the new ontology. The algorithm is highly modular as different strategies for handling higher-multiplicity mappings can be implemented and different (or multiple) matchers can be used. Our post-processing steps are based on the notion of conservativity but differently than what is usually done in the literature [13, 6, 19] we propose to remove axioms from the new ontology in order to repair many of the violations. This is important in order to keep ambiguity low and not have many classes with overlapping labels. We have formalised our framework, designed an exact general algorithm and also presented concrete approximate and practical algorithms. These have already been used in Babylon Health to build a medical Knowledge Base (using SNOMED, NCI, FMA, and CHV) that forms the data and knowledge backbone of various clinical services. Finally, we have conducted an experimental evaluation comparing our conservativity repairing approach to state-of-the-art mapping repair systems obtaining very encouraging results. In summary, our results verify that ambiguity is very-low (almost none introduced compared to the initial ambiguity of the input ontologies), there were no detectable violations (LDiff), no cycles, and our algorithm scales.

## References

1. Alchourrón, C.E., Gärdenfors, P., Makinson, D.: On the logic of theory change: Partial meet contraction and revision functions. *Journal of Symbolic Logic* 50(2), 510–530 (1985)
2. Ernst, P., Siu, A., Weikum, G.: KnowLife: a versatile approach for constructing a large knowledge graph for biomedical sciences. *BMC Bioinformatics* 16, 157:1–157:13 (2015)
3. Faria, D., Pesquita, C., Mott, I., Martins, C., Couto, F.M., Cruz, I.F.: Tackling the challenges of matching biomedical ontologies. *Journal of Biomedical Semantics* 9(1), 4:1–4:19 (2018)
4. Faria, D., Pesquita, C., Santos, E., Palmonari, M., Cruz, I.F., Couto, F.M.: The agreementmakerlight ontology matching system. In: *On the Move to Meaningful Internet Systems: OTM 2013 Conferences Confederated International Conferences: CoopIS, DOA-Trusted Cloud and ODBASE*. pp. 527–541 (2013)
5. Golbeck, J., Fragoso, G., Hartel, F.W., Hendler, J.A., Oberthaler, J., Parsia, B.: The national cancer institute’s thesaurus and ontology. *Journal of Web Semantics* 1(1), 75–80 (2003)
6. Jiménez-Ruiz, E., Grau, B.C., Horrocks, I., Llavori, R.B.: Ontology integration using mappings: Towards getting the right logical consequences. In: *Proceedings of the 6th European Semantic Web Conference, (ESWC)*. pp. 173–187 (2009)

7. Jiménez-Ruiz, E., Grau, B.C., Zhou, Y.: Logmap 2.0: towards logic-based, scalable and interactive ontology matching. In: Proc. of the 4th International Workshop on Semantic Web Applications and Tools for the Life Sciences. pp. 45–46 (2011)
8. Jiménez-Ruiz, E., Meilicke, C., Grau, B.C., Horrocks, I.: Evaluating mapping repair systems with large biomedical ontologies. In: Proceedings of the 26th International Workshop on Description Logics. pp. 246–257 (2013)
9. Konev, B., Walther, D., Wolter, F.: The logical difference problem for description logic terminologies. In: Proceedings of the 4th International Joint Conference on Automated Reasoning, IJCAR. pp. 259–274 (2008)
10. Manning, C.D., Surdeanu, M., Bauer, J., Finkel, J.R., Bethard, S., McClosky, D.: The stanford corenlp natural language processing toolkit. In: Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics, ACL. pp. 55–60 (2014)
11. Meilicke, C., Stuckenschmidt, H.: Applying logical constraints to ontology matching. In: Proceedings of the 30th Annual German Conference on Advances in Artificial Intelligence (KI). pp. 99–113 (2007)
12. Meilicke, C., Stuckenschmidt, H.: An efficient method for computing alignment diagnoses. In: Proceedings of the 3rd International Conference on Web Reasoning and Rule Systems, RR. pp. 182–196 (2009)
13. Meilicke, C., Völker, J., Stuckenschmidt, H.: Learning disjointness for debugging mappings between lightweight ontologies. In: 16th International Conference on Knowledge Engineering: Practice and Patterns (EKAW). pp. 93–108 (2008)
14. Rappaport, N., Nativ, N., Stelzer, G., Twik, M., Guan-Golan, Y., Stein, T.I., Bahir, I., Belinky, F., Morrey, C.P., Safran, M., Lancet, D.: Malacards: an integrated compendium for diseases and their annotation. Database (2013)
15. Salvadores, M., Alexander, P.R., Musen, M.A., Noy, N.F.: Bioportal as a dataset of linked biomedical ontologies and terminologies in RDF. Semantic Web 4(3), 277–284 (2013)
16. Schriml, L.M., Arze, C., Nadendla, S., Chang, Y.W., Mazaitis, M., Felix, V., Feng, G., Kibbe, W.A.: Disease ontology: a backbone for disease semantic integration. Nucleic Acids Research 40(Database-Issue), 940–946 (2012)
17. Shi, L., Li, S., Yang, X., Qi, J., Pan, G., Zhou, B.: Semantic health knowledge graph: Semantic integration of heterogeneous medical knowledge and services. BioMed Research International 2017(Article ID 2858423) (2017)
18. Shvaiko, P., Euzenat, J.: Ontology matching: State of the art and future challenges. IEEE Transactions on Knowledge and Data Engineering 25(1), 158–176 (2013)
19. Solimando, A., Jiménez-Ruiz, E., Guerrini, G.: Detecting and correcting conservativity principle violations in ontology-to-ontology mappings. In: Proceedings of the 13th International Semantic Web Conference (ISWC). pp. 1–16 (2014)
20. Solimando, A., Jiménez-Ruiz, E., Guerrini, G.: A multi-strategy approach for detecting and correcting conservativity principle violations in ontology alignments. In: 11th International Workshop on OWL: Experiences and Directions (2014)
21. Stoilos, G., Stamou, G., Kollias, S.: A string metric for ontology alignment. In: Proceedings of the International Semantic Web Conference (ISWC). LNCS, vol. 3729, pp. 624–637. Springer-Verlag (2005)
22. Vandebussche, P., Atemezing, G., Poveda-Villalón, M., Vatan, B.: Linked open vocabularies (LOV): A gateway to reusable semantic vocabularies on the web. Semantic Web 8(3), 437–452 (2017)