



Semi-automatic Integration of Learned Ontologies into a Collaborative Framework

Title	Semi-automatic Integration of Learned Ontologies into a Collaborative Framework
Author(s)	Nováček, Vít; Handschuh, Siegfried
Publication Date	2007

Semi-automatic Integration of Learned Ontologies into a Collaborative Framework^{*}

Vít Nováček¹, Loredana Laera², and Siegfried Handschuh¹

¹Digital Enterprise Research Institute, National University of Ireland, Galway
IDA Business Park, Lower Dangan, Galway, Ireland

E-mail: `first_name.last_name@deri.org`

²Department of Computer Science

University of Liverpool, UK

E-mail: `lori@csc.liv.ac.uk`

Abstract. The paper presents a novel ontology lifecycle scenario that explicitly takes the dynamics and data-intensiveness of real world applications into account. Changing and growing knowledge is handled by semi-automatic incorporation of ontology learning results into a collaborative ontology development framework. This integration bases mainly on automatic negotiation of agreed alignments, inconsistency resolution, an ontology versioning system and support of natural language generation tools, which alleviate the end-user effort in the incorporation of new knowledge. The architecture of the respective framework and notes on its progressive implementation are presented.

1 Introduction and Motivation

Ontologies on the Semantic Web, especially in case of real world applications, are very likely subject to change given the dynamic nature of domain knowledge. Knowledge changes and evolves over time as experience accumulates – it is revised and augmented in the light of deeper understanding; new facts are getting known while some of the old ones need to be revised and/or retracted at the same time.

This holds especially for scientific domains – we have to incorporate newly discovered facts and possibly change the inappropriate old ones in the respective ontology as the scientific research evolves further. However, even virtually any industrial domain is dynamic – changes typically occur in product portfolios, personal structure or industrial processes, which can all be reflected by an ontology in a knowledge management policy.

In these domains, ontology construction is usually the result of collaboration (which involves cooperation among ontology engineers and domain experts)

^{*} This work has been supported by the EU IST 6th framework's Network of Excellence 'Knowledge Web' (FP6-507482) and partially by Academy of Sciences of the Czech Republic, 'Information Society' national research program, the grant AV 1ET100300419.

through a manual process of the extraction of the knowledge. However, it is not always feasible to process all the relevant data and extract the knowledge from them manually, since we might not have a sufficiently large committee of ontology engineers and/or dedicated experts at hand in order to process new data anytime it occurs. This implies a need for (partial) automation of ontology extraction and management processes in dynamic and data-intensive environments. This can be achieved by ontology learning [15]. Therefore, a lifecycle of an ontology development process apt for universal application in scientific and/or industrial domains should also support appropriate mechanisms for dealing with the large amounts of knowledge that are *dynamic* in nature.

While there has been a great deal of work on ontology learning for ontology construction, e.g. [2], as well as on collaborative ontology development [18], relatively little attention has been paid to the integration of both approaches within an ontology lifecycle scenario. In this paper, we introduce our framework for practical handling of dynamic and large data-sets in an ontology lifecycle, focusing particularly on dynamic integration of learned knowledge into collaboratively developed ontologies. One of the key elements supporting our integration is the ability to reach an agreement on the semantics of the terms used in these ontologies. Since the ontologies are created under different circumstances and conditions and thus might represent different perspectives over similar knowledge, the achievement of an agreement will necessarily come through a (partially automated) negotiation process.

The dynamic nature of knowledge is one of the most challenging problems in the current Semantic Web research. Here we provide a solution for dealing with dynamics in large scale, based on properly developed connection of ontology learning and dynamic collaborative development. We do not concentrate on formal specification of respective ontology integration operators, we focus rather on implementation of them, following certain practical requirements:

1. the ability to process new knowledge (resources) automatically whenever it appears and when it is inappropriate for humans to incorporate it;
2. the ability to automatically compare the new knowledge with a “*master*” ontology that is manually and collaboratively designed and select the new knowledge accordingly;
3. the ability to resolve inconsistencies between the new and current knowledge, possibly favouring the assertions from presumably more complex and precise master ontology against the learned ones;
4. the ability to automatically sort the new knowledge according to user-defined preferences and present it to them in a very simple way, thus further alleviating human efforts in the task of final incorporation of the knowledge.

On one hand, using the automatic methods, we are able to deal with large amounts of changing data. On the other hand, the final incorporation of new knowledge is to be decided by the human users, repairing possible errors and inappropriate findings of the automatic techniques. The key to success and applicability is to let machines do most of the tedious and time-consuming work and provide people with concise and simple suggestions on ontology integration.

The rest of the paper is organized as follows: Section 2 presents a brief discussion of related work. Section 3 gives an overview of our ontology lifecycle scenario and framework, whereas Section 4 presents the integration of manually designed and learned ontologies in more detail. Finally, in Section 5 we give a simple illustrative example of concrete usage of the our integration approach. Section 6 concludes the paper and sums up our future work.

2 Related Work

Recent overviews of the state-of-the-art in ontologies and related methodologies can be found in [9]. However, none of them offers a direct solution to the previously mentioned problems. *Methontology* [8] is a methodology developed for designing ontologies to serve as a base for extending it towards evolving ontologies. The ODESeW and WebODE suite [4] projects provide an infrastructure and tools for semantic application development/management, which is also in the process of being extended for networked and evolving ontologies. However, they focus rather on the application development part of the problem than on the ontology evolution parts.

The above projects have all focused on either a single part of ontology evolution, or on a rather abstract study of the knowledge management cycle. However, mechanisms that would provide a clue on how to incorporate the dynamics into the lifecycle are typically put off only by introduction of the version management, which we find insufficient. Moreover, the need for automatic methods of ontology acquisition in data-intensive environments is acknowledged, but the place of the automatic techniques is usually not distinguished in the dynamic lifecycle settings. The work [10] describes a way of machine-assisted refinement of automatically learnt ontologies. Our approach [17] offers a broader picture of how to deal with the dynamics in a general lifecycle scenario. In this paper we concentrate on the combination of ontology learning and manual (collaborative) development in dynamic settings.

3 DINO – A Dynamic Ontology Lifecycle Scenario

DINO is an abbreviation of three key elements of our ontology lifecycle scenario and framework – *Dynamics*, *INtegration* and *Ontology*. However, the first two can also be *Data* and *INtensive*. All these features express the primary aim of our efforts – *to make the knowledge efficiently and reasonably manageable in data-intensive and dynamic domains*.

Figure 1 below depicts the scheme of the proposed dynamic and application-oriented ontology lifecycle that deals with the problems mentioned in the previous sections.

Our ontology lifecycle builds on four basic phases of an ontology lifecycle: *creation* (comprises both manual and automatic ontology development and update approaches), *versioning*, *evaluation* and *negotiation* (comprises ontology




Fig. 1. Dynamics in the ontology lifecycle

alignment and merging as well as negotiation among different possible alignments). The four main phases are indicated by the boxes annotated by respective names. Ontologies or their instances in time are represented by circles, with arrows expressing various kinds of information flow. The A boxes present actors (institutions, companies, research teams etc.) involved in ontology development, where A_1 is zoomed-in in order to show the lifecycle's components in detail.

The general dynamics of the lifecycle goes as follows. The community experts (or dedicated ontology engineers) develop a (relatively precise and complex) domain ontology (the *Community* part of the *Creation* component). They use means for continuous ontology *evaluation* and *versioning* to maintain high quality and manage changes during the development process. If the amount of data suitable for knowledge extraction is too large to be managed by the community, *ontology learning* takes its place. Its results are *evaluated* and partially (we take only the results with quality above a certain threshold into account) integrated into the more precise (but typically relatively small) reference community ontology. The integration is based on alignment and merging principles covered by the *negotiation* component. Its proposal and implementation principles form the key contribution of this paper (see Section 4 for details). The *negotiation* component takes its place also when interchanging or sharing the knowledge with other independent actors in the field. All the phases support ontologies in the standard OWL format [1], (namely in its OWL DL flavour), although some phases may not support the full range of the syntactic structures available (e.g. natural language generation from triples, as introduced in Section 4.7 here and in [16]). In the following we concentrate on the integration component. More information on other parts of the lifecycle can be found in [17].

4 Dynamic Integration of Newly Learned Knowledge in the DINO Framework

The key novelty of the presented lifecycle scenario is its support for incorporation of changing knowledge in data-intensive domains. This is achieved by implementation of a specific integration mechanism. Its scheme is depicted in Figure 2¹. The particular components and their connections are described in the following paragraphs.

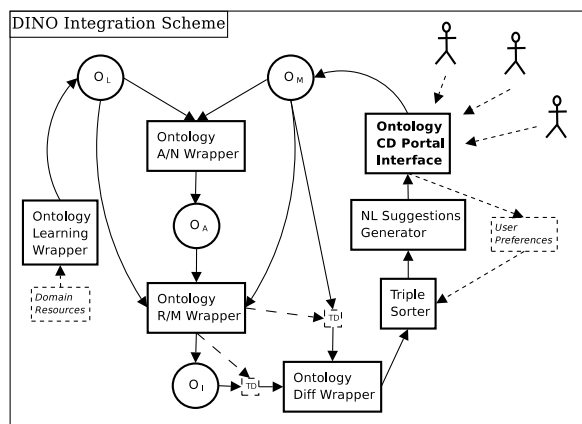


Fig. 2. Dynamic integration scheme

The integration scheme details the usage of generic lifecycle's components – mainly the *negotiation* and *versioning* – in the process of incorporation of learned ontologies into the collaboratively developed ones. However, the generic components serve only as a base for specific wrappers. Each of the phases of integration and their connections are described in the following sections.

4.1 Ontology Learning Wrapper

In this phase, machine learning and NLP methods are used for the processing of relevant resources and extracting knowledge from them (ontology learning). The ontology learning is realised using the Text2Onto framework [3]. We interface the toolbox indirectly within the collaborative ontology development portal based on MarcOnt Portal architecture (see Section 4.2). Configuration of the learning algorithms is set using a special user interface in the portal. The settings is used for batch processing of the new resources fed to the ontology learning component.

¹ The full arrows symbolise data flow in the scheme. Long-dashed arrows indicate production of a triple representation/dump of an ontology (the *TD* squares). The short-dashed arrows represent user involvement either in ontology production or extension preference sets' definition.

The results of one round of ontology learning – the O_L circle in Figure 2 – are optionally evaluated or refined using the Text2Onto confidence values and passed to the alignment/negotiation wrapper (see Section 4.3).

4.2 Ontology Collaborative Development Portal

The whole integration as well as the DINO framework is based on the MarcOnt Portal architecture [11] for collaborative ontology development. It is a part of a broader initiative aimed mainly at utilisation of various digital library development and maintenance efforts².

MarcOnt Portal offers domain-independent means for efficient distributed and collaborative ontology development. It supports features like ontology editing, ontology versioning (supported by the SemVersion system [21]), voting on ontology changes and evaluation of these votes. The elements of DINO realising various parts of the lifecycle are being implemented into the portal’s core, with access provided by respective new parts of portal’s user and administrative interfaces.

The ontology being developed using the portal’s collaborative interfaces is the *master reference ontology* in the whole lifecycle. It is also the source for the deployment of an official version of the ontology. The O_M circle in Figure 2 represents its dump that serves as a reference to be integrated with the O_L ontology resulting from the learning process. The final suggestions (see Section 4.7) form a base for a next version of the O_M ontology submitted after the integration.

4.3 Ontology Alignment/Negotiation (A/N) Wrapper

Once the learned ontology O_L and the master ontology O_M have been created, they need to be reconciled since they cover the same domain, but might be structured differently. The reconciliation of these ontologies depends on the ability to reach an agreement on the semantics of the terms used. The agreement takes the form of an alignment between the ontologies, that is, a set of correspondences (or mappings) between the concepts, individuals and properties in the ontologies. However, the ontologies are developed in different contexts and under different conditions and thus they might represent different perspectives over similar knowledge, so the process by which to come to an agreement will necessarily only come through a negotiation process. The negotiation process is performed using argumentation-based negotiation that uses preferences over the types of correspondences in order to choose the mappings that will be used to finally merge the ontologies (see Section 4.4). The preferences depend on the context and situation. A major feature of this context is the ontology, and the structural features thereof, such as the depth of the subclass hierarchy and branching factor, ratio of properties to concepts, etc. The analysis of the components of the ontology is aligned with the approach to ontology evaluation, demonstrated in [5], and can be formalized in terms of feature metrics. Thus

² See <http://www.marcont.org> for details on the whole MarcOnt Initiative.

the preferences can be determined on the characteristics of the ontology. For example, we can select a preference for terminological mapping if the ontology is lacking in structure, or prefer extensional mapping if the ontology is rich in instances.

Thus, the alignment/negotiation wrapper interfaces two tools – one for the ontology alignment discovery and one for negotiation of agreed alignment. We call these tools *AKit* and *NKit*, respectively, within this section. For the former, we use the ontology alignment API [6] developed by INRIA Rhone-Alpes³. For the negotiation we use the framework described in [13]. Both tools are used by the wrapper in order to produce O_A – an ontology consisting of axioms⁴ merging classes, individuals and properties in the O_L and O_M ontologies. It is used in consequent factual merging and refinement in the ontology reasoning and management wrapper (see Section 4.4 for details).

The wrapper itself works according to the meta-code in Algorithm 1. The

Algorithm 1 Meta-algorithm of the alignment and negotiation

Require: O_L, O_M — ontologies in OWL format
Require: *AKit*, *NKit* — ontology alignment and alignment negotiation tools, respectively
Require: *ALMSET* — a set of the alignment methods to be used
Require: *PREFSET* — a set of alignment formal preferences corresponding to the O_L, O_M ontologies (to be used in N-kit)

- 1: $S_A \leftarrow \emptyset$
- 2: **for** $method \in ALMSET$ **do**
- 3: $S_A \leftarrow S_A \cup AKit.getAlignment(O_L, O_M, method)$
- 4: **end for**
- 5: $A_{agreed} \leftarrow NKit.negotiateAlignment(S_A, PREFSET)$
- 6: $O_A \leftarrow AKit.produceBridgeAxioms(A_{agreed})$
- 7: **return** O_A

ontology alignment API offers several possibilities of actual alignment methods, which range from trivial lexical equality detection through more sophisticated string and edit-distance based algorithms to an iterative structural alignment by the OLA algorithm [7]. The ontology alignment API has recently been extended by a method for the calculation of a similarity metric between ontology entities, an adaptation of the SRMetric used in [20]. We also consider a set of justifications, that explain why the mappings have been generated. This information forms the basis for the negotiation framework that dynamically generates arguments, supplies the reasons for the mapping choices and negotiate an agreed alignment for both ontologies O_L and O_M .

4.4 Ontology Reasoning/Management (R/M) Wrapper

This wrapper is used for merging of the O_L and O_M ontologies. It uses Jena 2 Ontology API⁵. It merges the O_L and O_M ontologies according to the statements

³ See <http://alignapi.gforge.inria.fr/> for up-to-date information on the API.

⁴ Using constructs like *owl:equivalentClass*, *owl:sameAs*, *owl:equivalentProperty*, *rdfs:subClassOf* or *rdfs:subPropertyOf*.

⁵ See <http://jena.sourceforge.net/ontology/index.html>.

in O_A , preferring the lexical labels from the master O_M ontology when two labels are said to be equivalent. Moreover, the wrapper checks for (some) possible inconsistencies caused by the merging (using Jena’s simple OWL DL reasoner) and attempts to resolve them favouring the assertions in the O_M ontology, which are supposed to be more relevant, again. The resulting ontology O_I is passed to the ontology diff wrapper. As the Jena ontology model is internally based on a graph/triple (RDF) structure, it allows to easily export or transform an ontology in a triple format needed for the consequent wrapper (see Section 4.5 for details).

Algorithm 2 describes the meta-code of the process arranged by the ontology merging and reasoning wrapper. The inconsistency resolution is somewhat tricky.

Algorithm 2 Meta-algorithm of the merging and inconsistency resolution

Require: O_L, O_M, O_A — ontologies in OWL format
Require: $getEq()$ — function selecting all assertions of type *owl:equivalentClass*, *owl:sameAs*, *owl:equivalentProperty*
Require: $getRM()$ — function returning wrapper combining a generic ontology manager and (incomplete OWL Full) reasoner bound to the given ontology

- 1: $O_{tmp} \leftarrow copy(O_L)$
- 2: $O_I \leftarrow copy(O_M)$
- 3: $R_M \leftarrow getRM(O_M)$
- 4: $R_{tmp} \leftarrow getRM(O_{tmp})$
- 5: $R_L \leftarrow getRM(O_L)$
- 6: $R_A \leftarrow getRM(O_A)$
- 7: $R_I \leftarrow getRM(O_I)$
- 8: $equivalencies \leftarrow \{owl : equivalentClass, owl : sameAs, owl : equivalentProperty\}$
- 9: $UNIFIED \leftarrow \emptyset$
- 10: **for** $id \in getEq(O_A)$ **do**
- 11: $R_{tmp}.replaceLabels(id.O_L, id.O_M)$
- 12: $UNIFIED \leftarrow UNIFIED \cup id.O_M$
- 13: **end for**
- 14: $R_{ref} \leftarrow copy(R_{tmp})$
- 15: **for** $eq \in R_{tmp}.getAxiomsWithLabels(UNIFIED)$ **do**
- 16: $R_{tmp}.retractAxioms(eq)$
- 17: $R_I.addAxioms(eq)$
- 18: **end for**
- 19: $R_A.removeAxiomsOfType(equivalencies)$
- 20: $R_I.addAxioms(R_{tmp}.getAllAxioms())$
- 21: $R_I.addAxioms(R_A.getAllAxioms())$
- 22: $R_I.resolveInconsistencies(R_{ref})$
- 23: $R_I.augmentStructure()$
- 24: **return** O_I

However, we can apply a sort of “greedy” heuristic, considering the assertions in the master O_M ontology to be more valid. Therefore we query the R_{ref} structure (with the axioms of learned ontology, possibly with replaced labels) in the resolution process. We currently handle the following inconsistencies:

- **sub-class hierarchy cycles:** these are resolved by cutting the cycle by removing an *owl:subClassOf* statement present in R_{ref} ;
- **disjointness-subsumption** conflicts: if classes are said to be disjoint and a sub-class relationship holds between them or if they have common sub-classes at the same time, the conflicting assertion indicated by R_{ref} is removed;

- **disjointness-instantiation** conflicts: if an individual is said to be an instance of classes that are disjoint, the assertion indicated by R_{ref} is removed.

When there are several removal candidate axioms involved in one inconsistency, we sort them according to the confidence provided by the Text2Onto learning algorithms [3], which is stored in the R_{ref} reference structure. Similarly to [10], we start removing the axioms with least overall confidence, until we do not resolve the inconsistency (thus keeping the more “relevant” discoveries intact). We keep the conflicting assertions when they all originate from the O_M master ontology and let the users to cope with this fact. Note that the sources of inconsistencies are provided by simple natural language description and recorded for further examinations by human users – they can eventually decide to favour the learned assertions if appropriate for the given task in the given context.

The function *augmentStructure()* attempts to complete the structure of learned axioms using the more precise and complex knowledge in the O_M master ontology. Currently, augmentation of *owl:subClassOf* and instantiation relations using *rdfs:domain* and *rdfs:range* assertions in property definitions from O_M ontology is taken into account (see Section 5 for an example). More sophisticated extensions are possible in the future.

If we want to include even the “equal” labels from the learned ontology, we can omit the renaming and subtractions in lines 10-16 and 19 and include the respective equality statements from O_A into O_I , together with respective axioms from O_L . The decision depends on users – whether they want to prefer the labels from master ontology or not (e.g. when looking for possible unknown synonyms of important terms from O_M in domain resources; this could be useful for example in the medicine domain in task of identification of different names for the same drugs and/or proteins).

4.5 Ontology Diff Wrapper

Possible extension of a master ontology O_M by elements contained in the merged and refined ontology O_I naturally corresponds to the differences between them. These are discovered by means of the SemVersion library [21], which is interfaced within this wrapper. In particular, the possible extensions are equal to the additions O_I brings into O_M . We compute the additions from the triple-based representation⁶ of O_I and O_M ontologies. The additions are passed to the triple sorter then (see Section 4.6 for details).

4.6 Triple Sorter

The addition triples passed to this component form a base to the eventual extension suggestions for the domain experts. However, the number of additions

⁶ Since SemVersion does not currently support full OWL diff computations. The triple representation is provided by the ontology R/M wrapper, as indicated by the *TD* (triple dump) squares in Figure 2.

can generally be quite large, so an ordering that takes a relevance measure of possible suggestions into account is needed. Thus we can for example eliminate suggestions with low relevance level when presenting the final set to the users (without overwhelming them with a large number of possibly irrelevant suggestions).

As a possible solution to this task, we have proposed and implemented a method based on string subsumption and Levenshtein distance [14]. These two measures are used within relevance computation by comparing the predicate, subject and object lexical labels of a triple to two sets (S_p, S_n) of words, provided by users. The S_p and S_n sets contain preferred and unwanted words, respectively, concerning the lexical level of optimal extensions.

Relevance score of a triple T with respect to the wanted/unwanted sets of words S_p/S_n ($relScore(T, S_p, S_n)$ below) is given by the formula:

$$relScore(T, S_p, S_n) = rel(T, S_p) - rel(T, S_n),$$

where $rel(T, S)$ is a function measuring the relevance of the triple T with respect to the words in the set S . The higher the value, the more relevant the triple is. The function⁷ naturally measures the “closeness” of the T triple’s predicate, subject and object labels to the set of terms in S_w (S_w stands for S_p or S_n). The value of 1 is achieved when the label is a direct substring of or equal to any word in S_w or vice versa. When the Levenshtein distance between the label and a word in S_w is lower than or equal to the defined threshold t , the relevance decreases from 1 by a value proportional to the fraction of the distance and t . If this is not the case (i.e. the label’s distance is greater than t for each word in S_w), a similar principle is applied for words of the possibly multi-word label and the relevance is further proportionally decreased (the minimal possible value being 0).

4.7 Mapping Triples to Natural Language Suggestions

The DINO framework is supposed to be used primarily by users who are not experts in ontology engineering. Although the MarcOnt Portal [11] already offers a simple ontology editing interface, we would like to further help the user in ontology augmentation by the learned knowledge. Therefore the suggestions are produced in the form of very simple natural language statements. These are obtained directly from the sorted triples passed to this component, using a minor modification of the generation process in CLIE described in [19]. Examples of this final form of suggestions can be found in the next section. The suggestions are still bound to the underlying triples, therefore the user can very easily add the respective OWL axioms into the new version of the O_M master ontology without actually dealing with the intricate OWL syntax itself.

⁷ We described the relevance function in more detail in [17, 16], together with complexity analysis (which is in feasible class of $O(m \log(m))$ with respect to the number of triples).

5 Evaluation and Usage Example

The DINO framework is still a work in progress, and thus no proper evaluation has been carried out yet. However, preliminary evaluation of the core negotiation and preference-based suggestion sorting techniques has been made. The implemented sorting algorithm placed 80.7% of triples from a test sample into an order intuitively prepared by a human user. Details on the sorting evaluation are in [17, 16]. The negotiation component has been evaluated using the Ontology Alignment Evaluation Initiative test suite⁸. Experiments on the impact of the argumentation approach over a set of mappings and a comparison wrt. current alignment tools is presented in [12]. The preliminary results of these experiments are promising and suggest that the argumentation approach can be beneficial and an effective solution to the problem of dynamically aligning heterogeneous ontologies.

In the following we provide a simple illustrative example of concrete usage of the DINO integration mechanism. Imagine a medical institution that has developed an ontology O_M covering the basic concepts in clinical practice and research, possibly with help of ontology engineering experts when deploying the DINO framework. The ontology may need to be extended by new information in research (e.g. when new treatments or diagnosis methods are developed and published). Related information can be found in respective documents (research papers, industry white-papers, etc.). Figure 3 presents a sample text fragment with the respective learned OWL O_L ontology (we omit the namespace for simplicity).

```
... while cerebellar astrocytoma
is usually discovered by means of
CT... using a diagnostic procedure
of scanning... GVHD, an immune
dysfunction... GVHD, a disease being a type
of dysfunction...

...
<owl:ObjectProperty rdf:ID="discovered-by"/>
<owl:Thing rdf:ID="CT"/>
<owl:Thing rdf:ID="cerebellar-astrocytoma">
  <discovered-by rdf:resource="#CT"/>
</owl:Thing>
<owl:Class rdf:ID="diagnostic-procedure"/>
<owl:Class rdf:ID="immune-dysfunction"/>
<owl:Class rdf:ID="dysfunction"/>
<owl:Class rdf:ID="scanning">
  <rdfs:subClassOf rdf:resource="#diagnostic-procedure"/>
</owl:Class>
<immune-dysfunction rdf:ID="GVHD"/>
<owl:Class rdf:ID="disease">
  <rdfs:subClassOf rdf:resource="#dysfunction"/>
</owl:Class>
...
```

Fig. 3. A text sample and the learned ontology

The ontologies O_L and O_M are aligned and negotiated (see Figure 4). The preferences have been chosen on the basis of the ontological information of O_L and O_M (see Section 4.3 for details). The O_M ontology and the ontology O_A , consisting of axioms produced from the negotiated mappings are shown in Figure 5. When trying to merge the O_M and O_L ontologies into O_I according

⁸ See <http://oaei.ontologymatching.org/>.




Fig. 4. Negotiated mappings

```

...
<owl:ObjectProperty rdf:ID="InstrumentalProperty"/>
<owl:ObjectProperty rdf:ID="DiscoveredUsing">
  <rdfs:subPropertyOf rdf:resource="#InstrumentalProperty"/>
  <rdfs:range rdf:resource="#Manifestation"/>
  <rdfs:domain rdf:resource="#DiagnosisProcedure"/>
</owl:ObjectProperty>
<owl:Class rdf:ID="Manifestation"/>
<owl:Class rdf:ID="Procedure"/>
<owl:Class rdf:ID="DiagnosisProcedure">
  <rdfs:subClassOf rdf:resource="#Procedure"/>
</owl:Class>
<owl:Class rdf:ID="SoftTissueCytoma"/>
<owl:Class rdf:ID="AstroCytoma">
  <rdfs:subClassOf rdf:resource="#SoftTissueCytoma"/>
</owl:Class>
<owl:Class rdf:ID="Disease">
<owl:Class rdf:ID="Dysfunction">
  <rdfs:subClassOf rdf:resource="#Disease"/>
</owl:Class>
...
...
<owl:ObjectProperty rdf:ID="DiscoveredUsing">
  <owl:equivalentProperty rdf:resource="#discovered-by"/>
</owl:ObjectProperty>
<AstroCytoma rdf:ID="cerebellar-astrocytoma"/>
<owl:Class rdf:ID="DiagnosisProcedure">
  <owl:equivalentClass rdf:resource="#diagnostic-procedure"/>
</owl:Class>
<owl:Class rdf:ID="immune-dysfunction">
  <owl:subClassOf rdf:resource="#Dysfunction"/>
</owl:Class>
<owl:Class rdf:ID="Dysfunction">
  <owl:equivalentClass rdf:resource="#dysfunction"/>
</owl:Class>
...

```

Fig. 5. A master ontology sample and the respective mapping

to the technique described in Section 4.4, we find out that there is one inconsistency – “*disease*” is said to be a subclass of “*dysfunction*” and vice versa, which creates a cycle in the taxonomy. Therefore we remove the respective “invalid” assertion that originated from the O_L ontology. On the other hand, we can extend the learned knowledge based on range and domain of the “*DiscoveredUsing*” property. We can infer new assertions on the instantiation of “*cerebellar astrocytoma*” (instance of “*Manifestation*”) and “*CT*” (instance of “*DiagnosisProcedure*”).

Now we can produce the triples (with O_L equivalent labels replaced by those from O_M) from the O_I merge, together with respective suggestions based on the differences between O_I and O_M . We present the sorted triples and their transformations into natural language statements⁹ in Table 1.

Note that the above example may be also used if we just need to align and possibly extend the ontology with another institution’s knowledge base – the only difference is that we do not perform the ontology learning and also omit retractions in the integration process, as noted in Section 4.4. This can be applied in the critical task of inter-mediation of medicine information, for example.

⁹ They are preceded by sample relevance values, corresponding to {Scanning, discover, cytoma} and {subclass, disease, dysfunction} sets of preferred and unwanted labels, respectively.

<AstroCytoma rdf:ID="cerebellar-astrocytoma"/>	+0.667: CEREBELLAR ASTROCYTOMA is a <i>new instance</i> of ASTROCYTOMA.
<Manifestation rdf:ID="cerebellar-astrocytoma"/>	+0.667: CEREBELLAR ASTROCYTOMA is a <i>new instance</i> of MANIFESTATION.
<DiagnosisProcedure rdf:ID="CT"/>	+0.389: CT is a <i>new instance</i> of DIAGNOSIS PROCEDURE.
<immune-dysfunction rdf:ID="GVHD"/>	+0.333: GVHD is a <i>new instance</i> of IMMUNE DYSFUNCTION.
<owl:Class rdf:ID="scanning"> <rdfs:subClassOf rdf:resource="#DiagnosisProcedure"/> </owl:Class>	-0.444: A <i>new class</i> SCANNING is a <i>sub-class</i> of DIAGNOSIS PROCEDURE.
<owl:Thing rdf:ID="cerebellar-astrocytoma"> <discoveredUsing rdf:resource="#CT"/> </owl:Thing>	-0.667: CEREBELLAR ASTROCYTOMA is DISCOVERED USING CT.
<owl:Class rdf:ID="immune-dysfunction"> <rdfs:subClassOf rdf:resource="#Dysfunction"/> </owl:Class>	-0.833: A <i>new class</i> IMMUNE DYSFUNCTION is a <i>sub-class</i> of DYSFUNCTION.

Table 1. Extension triples and the respective NL suggestions

6 Conclusions and Future Work

We have presented the basic principles of DINO – a novel framework for dynamic ontology development in data-intensive domains. As a core contribution of the paper, we have described the mechanism of integration of learned and collaboratively developed knowledge. It covers all the requirements specified in Section 1. The proposed combination of automatic and collaborative tools in knowledge acquisition, integration and inconsistency resolution ensures production of reliable, broad and precise ontologies when using DINO in dynamic settings.

Our present and future work concentrates mainly on full implementation of the DINO framework by the respective extensions of the MarcOnt Portal architecture. We also plan to continuously evaluate the framework in line with demands of medicine industry (which we are currently specifying with help of e-Health experts) and possibly other domains.

References

1. S. Bechhofer, F. van Harmelen, J. Hendler, I. Horrocks, D. L. McGuinness, P. F. Patel-Schneider, and L. A. Stein. *OWL Web Ontology Language Reference*, 2004. Available at (February 2006): <http://www.w3.org/TR/owl-ref/>.
2. C. Brewster, F. Ciravegna, and Y. Wilks. User-centred ontology learning for knowledge management. In *In Proceedings 7th International Workshop on Applications of Natural Language to Information Systems, Stockholm.*, 2002.
3. P. Cimiano and J. Völker. Text2Onto - a framework for ontology learning and data-driven change discovery. In *Proceedings of the NLDB 2005 Conference*, pages 227–238. Springer-Verlag, 2005.
4. O. Corcho, A. Lopez-Cima, and A. Gomez-Perez. The ODESeW 2.0 semantic web application framework. In *Proceedings of WWW 2006*, pages 1049–1050, New York, 2006. ACM Press.
5. K. Dellschaft and S. Staab. On how to perform a gold standard based evaluation of ontology learning. In *Proceedings of the International Semantic Web Conference. Athens, GA, USA.*, 2006.
6. J. Euzenat. An API for ontology alignment. In *ISWC 2004: Third International Semantic Web Conference. Proceedings*, pages 698–712. Springer-Verlag, 2004.
7. J. Euzenat, D. Loup, M. Touzani, and P. Valtchev. Ontology alignment with ola. In *Proceedings of the 3rd International Workshop on Evaluation of Ontology based Tools (EON)*, Hiroshima, Japan, 2004. CEUR-WS.

8. M. Fernandez-Lopez, A. Gomez-Perez, and N. Juristo. Methontology: from ontological art towards ontological engineering. In *Proceedings of the AAAI97 Spring Symposium Series on Ontological Engineering*, pages 33–40, Stanford, USA, March 1997.
9. A. Gomez-Perez, M. Fernandez-Lopez, and O. Corcho. *Ontological Engineering*. Advanced Information and Knowledge Processing. Springer-Verlag, 2004.
10. P. Haase and J. Völker. Ontology learning and reasoning - dealing with uncertainty and inconsistency. In P. C. G. da Costa, K. B. Laskey, K. J. Laskey, and M. Pool, editors, *Proceedings of the Workshop on Uncertainty Reasoning for the Semantic Web (URSW)*, pages 45–55, NOV 2005.
11. S. Kruk, J. Breslin, and S. Decker. MarcOnt initiative. Lión Deliverable 3.01, DERI, Galway, 2005.
12. L. Laera, I. Blacoe, V. Tamma, T. Payne, J. Euzenat, and T. Bench-Capon. Argumentation over ontology correspondences in mas. In *In Proceedings of the Sixth International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS 2007), Honolulu, Hawaii, USA. To Appear*, 2007.
13. L. Laera, V. Tamma, J. Euzenat, T. Bench-Capon, and T. R. Payne. Reaching agreement over ontology alignments. In *Proceedings of 5th International Semantic Web Conference (ISWC 2006)*. Springer-Verlag, 2006.
14. V. I. Levenshtein. Binary codes capable of correcting deletions, insertions and reversals. *Cybernetics Control Theory*, 10:707–710, 1966.
15. A. Maedche and S. Staab. Ontology learning. *Handbook on Ontologies*, 2004.
16. V. Nováček, M. Dabrowski, S. R. Kruk, and S. Handschuh. Extending community ontology using automatically generated suggestions. In *Proceedings of FLAIRS 2007*. AAAI Press, 2007. In press.
17. V. Nováček, S. Handschuh, L. Laera, D. Maynard, M. Völkel, T. Groza, V. Tamma, and S. R. Kruk. Report and prototype of dynamics in the ontology lifecycle (D2.3.8v1). Deliverable 238v1, Knowledge Web, 2006.
18. Y. Sure, M. Erdmann, J. Angele, S. Staab, R. Studer, and D. Wenke. OntoEdit: Collaborative Ontology Development for the Semantic Web. In *1st International Semantic Web Conference (ISWC2002)*, Sardinia, 2002. Springer.
19. V. Tablan, T. Polajnar, H. Cunningham, and K. Bontcheva. User-friendly ontology authoring using a controlled language. In *Proceedings of LREC 2006 - 5th International Conference on Language Resources and Evaluation*. ELRA/ELDA Paris, 2006.
20. V. Tamma, I. Blacoe, B. L. Smith, and M. Wooldridge. Introducing autonomic behaviour in semantic web agents. In *In Proceedings of the Fourth International Semantic Web Conference (ISWC 2005), Galway, Ireland, November.*, 2005.
21. M. Völkel and T. Groza. SemVersion: RDF-based ontology versioning system. In *Proceedings of the IADIS International Conference WWW/Internet 2006 (ICWI 2006)*, 2006.