CleanONTO: Evaluating Taxonomic Relationships in Ontologies

Derek Sleeman
Department of Computing Science
University of Aberdeen
Aberdeen, AB24 3FX
sleeman@csd.abdn.ac.uk

Quentin Reul
Department of Computing Science
University of Aberdeen
Aberdeen, AB24 3FX
qreul@csd.abdn.ac.uk

ABSTRACT

Consistent ontologies are vital for the growth of the Semantic Web. We describe and appraise the OntoClean methodology and the different implementations available to evaluate taxonomic relationships in ontologies. We propose a new system, CleanONTO, which uses definitions to describe each concept, where definitions are paths from the concept to the root node of the ontology. In the current study, these definitions (paths) have been extracted from WordNet.

1. INTRODUCTION

The use of ontologies in applications is progressively more common. An ontology has been defined by Noy & McGuiness [12] as: *"a common vocabulary for researchers who need to share information in a domain. It includes machine-interpretable definitions of basic concepts in the domain and relations among them."*

The sharing of information requires ontologies to be consistent and therefore the evaluation of ontologies is an important activity. Several methodologies have been suggested over the years. On the one hand, qualitative approaches evaluate an ontology by asking the user to rate an ontology, or subparts of an ontology. Gomez-Perez [5] presents an evaluation based on different criteria (consistency, completeness, conciseness, expandability, and sensitiveness). Welty & Guarino [15][7][8] have argued that many ontologies have inconsistent taxonomic relationships. As part of the OntoClean methodology, they introduce a coding scheme for concepts which uses basic notions, such as *essence, identity, unity,* and *dependence.* On the other hand, Brewster et al. [1] suggest one decomposes the ontology into concepts and relations to evaluate its *'goodness of fit'* for particular sets of natural language texts, a corpus; this evaluation is done using statistical approaches. However, CleanONTO uses definitions to describe each concept, where definitions are paths from the concept to the root node of the ontology. In the current study, these definitions (paths) have been extracted from WordNet 2.1 [11][2] rather than a domain specific corpus.

As part of this paper, we describe and then criticize the methodology used as part of OntoClean (sections 2 & 3). In section 4, we briefly describe WordNet before proposing a new approach to evaluate ontologies using concept definitions (paths) which are currently derived from WordNet; we refer to this system as CleanONTO. In section 6, we report the main differences between OntoClean and our approach. We then outline planned future work, and finally we discuss how CleanONTO fits within the ontology management framework being developed at Aberdeen.

WWW2006, May 22–26, 2006, Edinburgh, UK.
.

2. ONTOCLEAN METHODOLOGY

The OntoClean approach focuses on the arguments (properties or concepts) involved in taxonomic relationship. The semantics of subsumption (or IS-A relation) is the standard one: if P and Q are unary predicate symbols,

DEFINITION 1. *P IS-A Q iff* $I(P) \subseteq I(Q)$

where I is an interpretation function mapping unary predicates into subsets of the domain [6]. In order to solve subsumption inconsistencies, Welty & Guarino [15][8] introduce meta-properties and show how they impose some constraints.

To explain their approach, they used an "unclean" ontology (Figure 1 from [8]). The ontology has been created using taxonomic pairs from WordNet[1], and Pangloss[2].

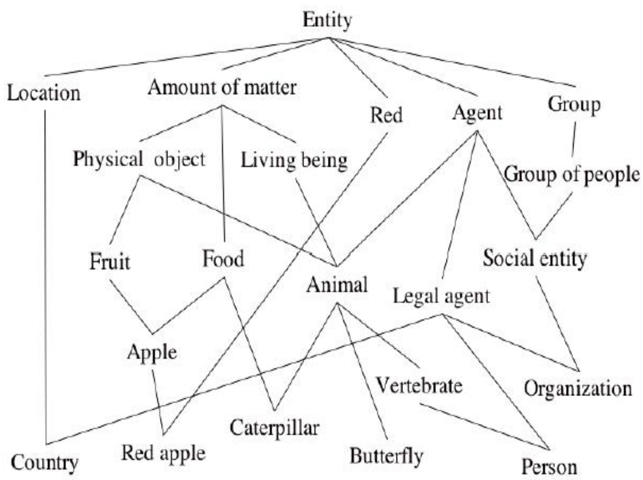

Figure 1: Unclean Ontology

As part of the methodology they introduced a set of meta-properties drawn from philosophy, namely *unity*, *identity*,

[1]http://wordnet.princeton.edu/
[2]http://www.lti.cs.cmu.edu/Research/Pangloss/

essence and *dependence*, which are used to characterize relevant aspects of the intended interpretation of the properties, classes, and relations that make up an ontology. Moreover, the meta-properties dictate several constraints on the taxonomic structure of an ontology, which can be used to evaluate an ontology. The following sections describe each meta-property in some detail.

2.1 Unity

The notion of unity is related to the problem of recognizing all the parts that forms an instance from the rest of the world by means of the *unifying relation*. It is useful to determine the intended meaning of properties in an ontology. Certain properties have wholes (an entity lacking no part) as instances and others have not. For example, the link between *'Food'* and *'Amount of matter'* is deemed to be consistent as neither have wholes as instances.

In addition, Welty & Guarino define *unity criteria* (UC) as: *"the specific conditions that must hold among the parts of a certain entity in order to consider it whole"*. Using meta-properties, they distinguish: *"properties all whose instances must carry a common UC from those that do not. Among the latter, we further distinguish properties all of whose instances must be wholes, although with different UCs, from properties all of whose instances are not necessarily wholes"*.

Using the above definitions on the taxonomy, we say that *'Animal'* carries unity (**+U**), due to being subsumed by *'Living being'* (**+U**), *'Legal Agent'* carries no unity (**-U**), and *'Amount of matter'* carries anti-unity (~**U**). With these meta-properties assigned, we can determine whether links are consistent. For example, the relation between *'Amount of matter'* and *'Living being'* is considered to be inconsistent as a ~U property can not subsume one with +U.

2.2 Identity

Identity is based on intuitions about how people interact with individual entities in the world around them. It is concerned with the problem of distinguishing a specific instance of a certain class from other instances of this class by means of analyzing essential characteristic property, called *identity criteria*. These properties may include the constituent of a certain entity (e.g. the chemical constitution of acid), the topological relationships among constituents (e.g. arrangement of nodes in a graph), or the persistence of a certain property over time (e.g. the permanence of a certain status).

OntoClean [6] proposes to distinguish between properties that carry an identity criterion (**+I**) and those that do not (**-I**). Moreover, it makes a further distinction to mark properties that supply their "own" identity criteria (**+O**), which are not inherited from the subsuming properties.

Looking at Figure 1, we find a subsumption relation between *'Physical object'* and *'Animal'*. However, this relation is deemed to be inconsistent. *Being alive* is an essential property for animals, as they cease to exist at death, whereas it is not essential to physical bodies. Under these assumptions, an animal could be at the same time necessarily alive and not necessarily alive suggesting the need for two different entities and the subsumption link should be removed. In the "unclean" taxonomy, we can also find a IS-A link between *'Living being'* and *'Animal'*. This relation is considered consistent as they both carry the same identity criteria (being alive).

2.3 Essence

In the case of identity, we need to accept that an individual may remain the same while displaying different properties at different times. This problem leads to the notion of an *essential property*, on which the definition of rigidity is based.

A property is considered rigid (**+R**) if it is essential to all its possible instances, which means that it cannot stop being an instance of that property in a different domain. For example, *'Physical object'* is +R as we assume that every instance ceases to exist when it ceases to be a physical object. Furthermore, there are non-rigid (**-R**) properties, which can gain or lose their instances. For instance, *'Red'* is judged to be -R since some instances of it may be necessarily so, and most will not. A further distinction is made among non-rigid properties, namely the introduction of anti-rigid (~**R**) properties, which are not essential to all their instances. *'Food'* is considered to be ~R as anything that is food can also not be food.

Using the above definition, the expert is able to analyze the links as part of the taxonomy (Figure 1). Let's look at the link between *'Food'* and *'Apple'*. We have already determined the rigidity property for *'Food'*. The taxonomy defines *'Apple'* to be indirectly subsumed from *'Physical object'* as an apple ceases to exist once it has been eaten, and therefore we assign +R to it. Using these meta-properties and the premise that ~R properties can not subsume +R properties, the link between the two entities has to be removed.

2.4 Dependence

The fourth notion, *ontological dependence*, involves every instance dependent on a concept to exist. Another way to view dependency is to differentiate between *intrinsic* and *extrinsic* properties. An intrinsic property is typically independent, such as *'Food'*. Whereas extrinsic properties have a relational nature with other instances, such as *'Apple'*.

2.5 Cleaned taxonomy

The cleaned taxonomy resulting from applying the OntoClean methodology is shown in Figure 2 from [8]. Note that these results were obtained with the enhanced OntoClean, namely a version of the system which had addressed the various points raised by Kaplan [9]. This taxonomy has fewer "multiple inheritance" links than the original. This is due to the misuse of subsumption; subsumption is sometimes used to represent other relations. For example, we could say that an animal is a physical object, however we have shown in section 2.2 that this is not logically consistent with the subsumption relationship.

3. ONTOCLEAN APPRAISAL

Although we believe that the OntoClean approach is addressing a very real problem, we are not convinced that the meta-properties schema introduced is usable by knowledge engineers or domain experts, as different knowledge engineers tend to describe the same concept with significantly different sets of meta-properties.

In this section, we describe some of the work that has been done in ontology evaluation which has appraised OntoClean's approach. Some of the approaches also suggest alternatives.

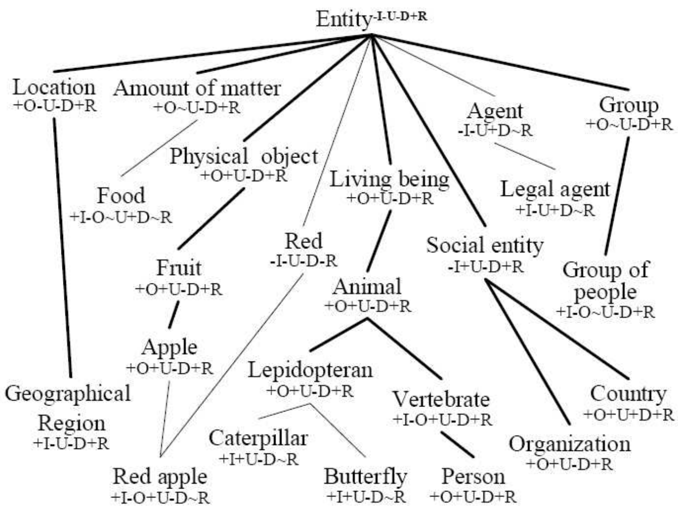

Figure 2: The final cleaned taxonomy as produced by OntoClean

3.1 Automatic Evaluation of Ontologies

An important study by Volker et al. [14] reports that fairly experienced knowledge engineers were only 38% consistent in using the four meta-properties provided by OntoClean; dependency seems to be the least reliably reported.

The paper then describes a very intriguing approach where they define patterns associated with each of the meta-properties, and then use these patterns to search the web for instances (of the meta-properties). An example of a pattern to search for phrases such as "is no longer..." in order to determine **Rigidity** is shown in Figure 3.

```
<pattern>
  <variable name="x" />
  <evidence type="false" for="R" />
  <google regex="is\t\w+ no\t\w+ longer\t(DT\w+\t)?(NN|NP
      |NNS|NPS) x\t[^(NN|NP|NNS|NPS)]">
    <query string="is no longer a x" />
    <query string="is no longer an x" />
    <query string="is no longer x" />
  </google>
</pattern> %
```

Figure 3: Pattern for Negative Evidence for Rigidity (R)

These results were then subjected to linguistic filtering, and this data was used to train classifiers for each of the meta-properties. Encouraging results are reported using these classifiers [14].

3.2 Data Driven Ontology Evaluation

Brewster et al. [1] argues for a clear set of evaluation methodologies and makes the case for an approach based on a data-driven evaluation methodology. Further, they define a good ontology as one that serve its purpose.

The authors identify problems of using **precision** (amount of knowledge correctly identified with respect to the whole knowledge available in the ontology) and **recall** (amount of knowledge correctly identified with respect to all the knowledge that it should have identified) to perform ontology evaluation; these techniques were originally developed in Information Retrieval and Natural Language Processing.

In order to solve the aforementioned problems, they state: *"an ontology is composed of concepts and relations, some of which are explicitly defined, others which follow from a set of axioms"*. The approach suggests finding "signatures" in natural language texts for relevant concepts (from the ontology) by identifying salient terms used in a corpus of texts for a given domain. They propose a **Probabilistic Approach** to determine the *"best fit"* between a given corpus and a set of ontologies. In their view, many difficulties in the matching process are caused by differences in the granularity of the sets of concepts being compared.

4. WORDNET

WordNet [11][2] is a lexical database, which arranges English words into sets of synonyms called synsets, provides short definitions, and records the various relations between these synonym sets.

As part of this work, we are mainly looking at the hyponymy/hypernymy relation, which inherits its entire feature from its superordinate with the addition of at least one concept. This is the basic organisation of nouns in WordNet.

5. CLEANONTO

We have recently developed an alternative ontology checking system, CleanONTO, which is based on having definitions for each of the concepts in the ontology/taxonomy. In our approach each of the concepts is defined as a path where the path for concept C contains itself, its parent, and its parent's parent up to the root node *'Entity'*. We give below the definitions (or paths) for the concepts *'Animal'* and *'Food'*.

- *Concept: Animal*
  animal → organism → living_being → physical_object → entity
- *Concept: Food*
  food → solid → matter → entity

For the present we have extracted these paths from WordNet 2.1 [2] as it is able to give us the sort of information that we require. However, we are well aware of the criticism which have been made of WordNet as a source of ontological information [4][13][3], but wish to stress that we are attempting to develop an approach to ontology evaluation which is based not specifically on WordNet but is a definition-based (path-based) approach. We discuss in future work plans to acquire these definitions from other sources.

5.1 Methodology

Our overall approach has three distinct phases. Firstly, each of the concepts which occurs in the "unclean" ontology was looked up in the online version of WordNet[3] by an investigator who then reported each of the *"paths"*. Secondly, we analyze the ontology and break all links which are inconsistent (Definition 2). Sometimes this results in an *"orphan"* node or subtree. This process also results in a basic tree, which we refer to as Tree-0.

DEFINITION 2. *A concept **B** subsumes a concept **A**, iff the whole path of **B** is found in the path of **A**, where concept **A** is called the child, and **B** is called the parent*

[3]http://wordnet.princeton.edu/perl/webwn

Table 1: Paths for concepts found in WordNet

Concepts	Paths
Entity	entity
Amount of matter	matter → entity
Food	food → solid → matter → entity
Fruit	fruit → reproductive_structure → plant_organ → plant_part → natural_object → physical_object → entity
Apple	apple → edible_fruit → fruit → reproductive_structure → plant_organ → plant_part → natural_object → physical_object → entity
Physical object	physical_object → entity
Living being	living_being → physical_object → entity
Animal	animal → organism → living_being → physical_object → entity
Person	person → organism → living_being → physical_object → entity
Agent	agent → representative → negociator → communicator → person → living_being → physical_object → entity
Group	group → abstraction → abstract_entity → entity
Organization	organization → social_group → group → abstraction → abstract_entity → entity

The third phase is to place these orphan nodes and subtrees back onto the tree so that the consistency rules are not violated. In fact there may be several nodes to which these orphans can be attached. The next subsections describe these processes in some more detail.

5.1.1 Acquiring descriptions for each concept

In phase one, one of the investigators looked up each of the concepts which occurred in the "unclean" ontology (Figure 1) using the interactive version of WordNet. This provides what we call a path from the named concept (e.g. *'Person'*) to the root node of the ontology, i.e. *'Entity'*. In some cases a concept has several senses and the investigator chose what he thought was the appropriate sense of the concept. Table 1 shows a selection of the concepts and their associated paths.

5.1.2 Breaking inappropriate links

There are two reasons why a link may be considered to be inconsistent. Firstly, if one of the nodes does not occur in the dictionary. For example, in the "unclean" taxonomy, there are three concepts, namely *'Group of people'*, *'Legal agent'*, *'Social entity'* and *'Red apple'*, which do not occur in the dictionary and hence the following links are broken:

- *'Legal agent'* & *'Agent'*
- *'Social entity'* & *'Agent'*
- *'Group of people'* & *'Group'*
- *'Social entity'* & *'Group'*
- *'Person'* & *'Legal agent'*
- *'Organization'* & *'Legal agent'*
- *'Organization'* & *'Social entity'*
- *'Red apple'* & *'Red'*
- *'Red apple'* & *'Apple'*

Secondly, CleanONTO inspects each remaining link in the ontology to check if it conforms to Definition 2, if not the link is broken. Below, we present examples of consistent relationships, as well as inconsistent ones from the "unclean" taxonomy (Figure 1).

Examining the relation between *'Amount of matter'* and *'Food'*, paths for both these concepts are found in Table 1 and they are

- *Concept: Amount of matter*
  matter → entity
- *Concept: Food*
  food → solid → matter → entity

where *'Amount of matter'* is the parent and *'Food'* the child node. As all the elements of the parent path are found in the child path, we consider this link to be consistent. Note that in this particular subsumption analysis[4] we ignore additional and intermediary concepts, in this case the *'solid'* concept.

When we apply the same analysis to the link between *'Amount of matter'* and *'Living being'* we firstly retrieve the two concepts from the table:

- *Concept: Amount of matter*
  matter → entity
- *Concept: Living being*
  living_being → physical_object → entity

and here the matching algorithm concludes that the elements in the parent path are not contained in the child path, and hence the link is inconsistent.

As a result of applying this second process to all the remaining links in the taxonomy, we note that the following links are broken:

- *'Animal'* & *'Agent'*
- *'Animal'* & *'Physical object'*
- *'Apple'* & *'Food'*
- *'Caterpillar'* & *'Food'*

and we identified the following concepts as "orphans":

- *'Person'* (A)
- *'Organization'* (B)

Furthermore, the algorithm removes two subtrees from the ontology (Figures 4 & 5) leaving the basic ontology, Tree-0 (Figure 6).

[4]This is implemented through the matching algorithm.

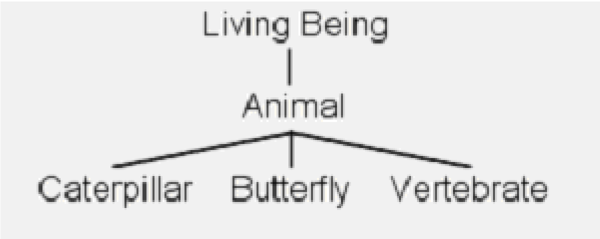

Figure 4: Subtree removed from ontology (C)

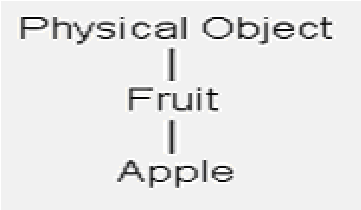

Figure 5: Subtree removed from ontology (D)

5.1.3 Creating a consistent tree

Once every inconsistent entities and sub-trees have been removed, we try to place these elements back onto Tree-0 (Figure 6), so that all the links in the enhanced tree are consistent.

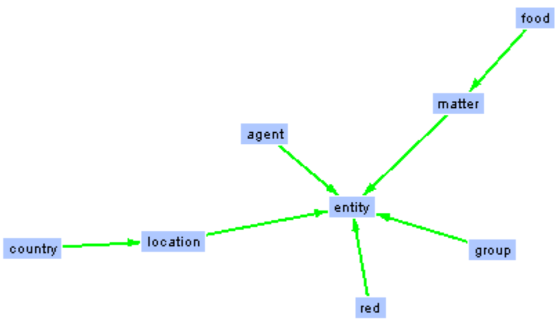

Figure 6: Resulting Tree 0 in CleanONTO

We first attempt to add the concept *'Person'* (A) to Tree-0. The only node which satisfies the matching criteria which we give above is the concept *'Entity'* and so *'Person'* is added directly to *'Entity'*. In making this connection we have effectively "ignored" the three intermediary concepts in the path for *'Person'*, namely *'Physical object'*, *'Living being'* and *'Organism'*. But we see below in the reappraisal step how some of these concepts are subsequently added to the evolving tree. The above step creates a new consistent tree (Tree-1).

Secondly, we examine where we could add *'Organization'* (B) on Tree-1. Looking at the paths found in Table 1, we observe that the node could be added to either *'Entity'* or *'Group'* as their paths are found in the child path. As more elements from the *'Group'* path are present in the *'Organization'* path, we add a link between those concepts forming Tree-2. A reappraisal of the previously added node (i.e. *'Person'*) is then performed. In this case no modification is made to the tree.

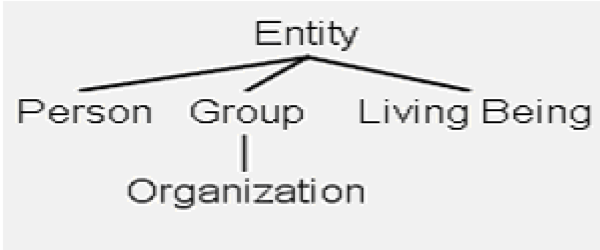

Figure 7: Tree 3.1

We then consider where to add the first subtree (Figure 4), by comparing its root node (*'Living being'*), against all the concepts in Tree-2. The algorithm decides that it is only acceptable to add *'Living being'* to *'Entity'*, which leads to Tree 3.1 (Figure 7).

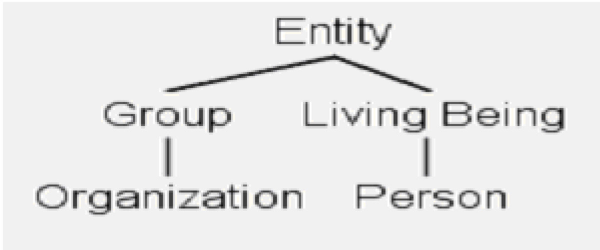

Figure 8: Tree 3.2

Reappraisal of *'Person'* is triggered by the appearance on the tree of an intermediary node in its path (i.e. organism or living_being or physical_object). In this instance, the re-arrangement is triggered by the addition of *'Living being'* (Figure 8).

Lastly, we determine where to incorporate the second subtree (Figure 5), by identifying where to place its root node (*'Physical object'*). The addition of the node *'Physical object'* to the *'Entity'* node causes a reappraisal of the tree as the *'Physical object'* concept also occurs as an intermediary concept in the *'Person'* concept. Once this rearrangement has been made, the resulting taxonomy is as shown in Figure 9.

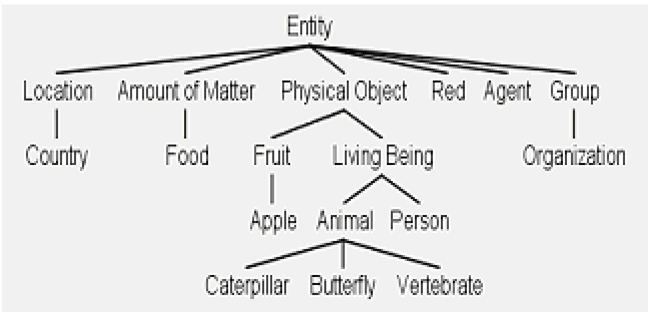

Figure 9: Cleaned ontology as produced by CleanONTO

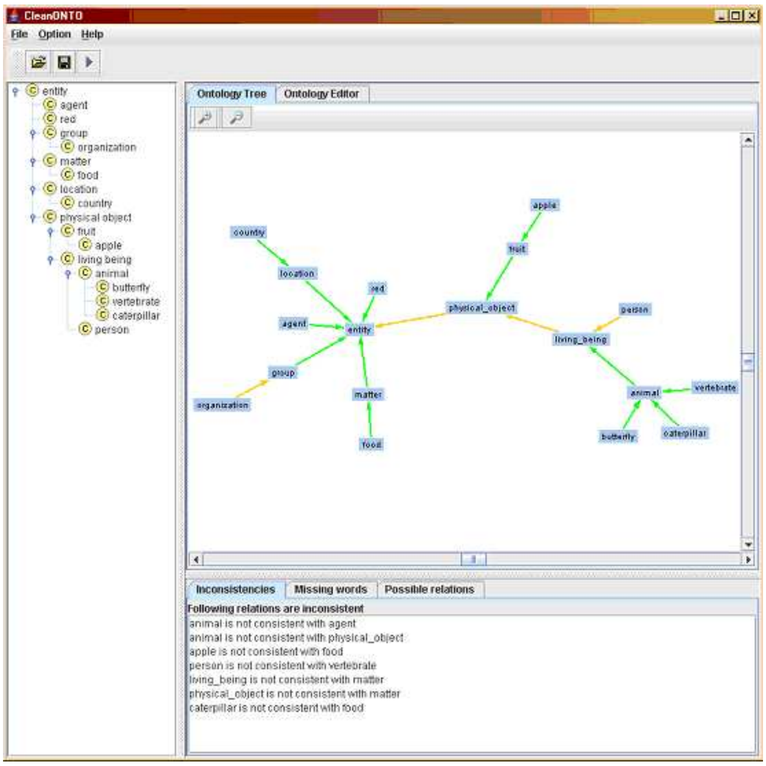

Figure 10: The CleanONTO GUI

5.2 Implementation

We have recently implemented the CleanONTO system (Figure 10), which when provided with an OWL[5] ontology, and a file produced manually using the WordNet online search[6] facility, automatically creates a consistent ontology. The application is written in Java and uses Jena API[7] to process the ontology.

The GUI contains three main areas. The left-hand side area shows the resulting ontology as a class tree, similar to the one found in *Protégé*. The upper main area contains two tabs. The first tab holds a graphical representation, using TouchGraph[8], of the resulting ontology. Only concepts found in WordNet are represented. The links that have been identified as consistent during the "Breaking inappropriate links" stage are represented in green. Links that are added in the reconstruction stage are represented in orange. The second tab allows the user to read the original ontology. However, the user is unable to make any changes to it. The lower area is subdivided in three different tabs. The first tab lists all inconsistent links found in the ontology. The second tab lists all concepts that were not found in WordNet. The last lists relations resulting from the "Creating a consistent tree" stage of the algorithm.

[5]http://www.w3.org/TR/owl-features/
[6]http://wordnet.princeton.edu/perl/webwn
[7]http://jena.sourceforge.net/ontology/index.html
[8]http://www.touchgraph.com/

6. DIFFERENCES BETWEEN ONTOCLEAN AND CLEANONTO

Comparing the ontologies produced by OntoClean and CleanONTO for this problem (Figures 2 & 9 respectively), we note that there are significant differences between them. Firstly, the concepts in the taxonomy which are not in WordNet are not in CleanONTO's revised tree; these concepts are listed in section 5.1.2.

Secondly, OntoClean makes distinctions based on meta-properties, whereas we compare their hierarchical "descriptions" extracted from WordNet. As explained in the OntoClean section, both concepts *'Physical object'* and *'Living being'* are deemed to have different properties and therefore the link between them is removed. Whereas, our approach finds the path of *'Physical object'* in the path of *'Living being'* and therefore accepts the link as consistent.

Thirdly, OntoClean sometimes adds a new concept to the taxonomy to show the different levels. For example, the concept *'Lepidopteran'* has been added to *'Animal'*. Whereas, this type of activity is proposed as part of our future work (section 7).

Finally, it is an interesting and important question as to which of the ontologies produced by OntoClean and CleanONTO, knowledge engineers find more acceptable, and **why**. We plan to include this comparison in the empirical evaluations to be carried out shortly, see section 7.

7. CONCLUSION & FUTURE WORK

In this paper, we have described an alternative coding scheme based on paths extracted from WordNet to evaluate ontologies. At present, CleanONTO removes all inconsistent relations from the taxonomic structure, and produces a consistent ontology.

Additionally, we have planned the following improvements to the system:

1. Use of WordNet API: We are looking at the possibility of using a WordNet API in order to automatically create the "paths" needed to evaluate an ontology. At present, we have created these manually which clearly limits the usefulness of the approach.

2. Optimization phase: At present only those nodes which violate the consistency rule are removed but it is possible (under some circumstances) that a more compact taxonomy could be produced by further rearranging "legal" nodes in the taxonomy. The algorithm must of course still ensure that all the resulting links are consistent.

3. Inclusion of intermediary nodes: The taxonomy construction algorithm discussed above is minimalist, in that intermediary nodes are only included if they appear on the path of two concepts (e.g. living_being in Figure 9). Some knowledge engineers might be unhappy with this strategy and so we plan to implement a further reconstruction option in which all intermediary nodes are added to the taxonomy.

4. Transparency of decisions: We plan to add an explanation facility to the system, so the user can ask questions such as why the link between A&B is inconsistent, or why the link between A&C is consistent.

5. Evaluation of the resulting ontology: We plan to show several knowledge engineers an example of an "unclean" ontology and ask them what changes they would make to the taxonomy and why. We then plan to show them the revised ontology produced by CleanONTO, and ask them to critique it, and to compare the revised ontology with the changes they suggested be made to the "unclean" ontology. In a second experiment, we plan to show the knowledge engineers the "unclean" ontology, ask them to suggest how it might be modified, and then show them the ontologies produced by OntoClean and CleanONTO. We then plan to ask them to critique both of the revised ontologies and to say which one they preferred and why. These empirical studies will be completed before the Workshop and the results will be presented.

6. Using an alternative to WordNet: The approach describe above compares the path of a parent concept with that of a child concept. To date, as discussed in section 5, we have acquired these paths from WordNet, but clearly the above algorithm is independent of the source of these paths. So if we find a source of better (e.g. conceptually more accurate) paths, this approach could be adapted with minimal effort.

8. ABERDEEN'S ONTOLOGY MANAGEMENT ENVIRONMENT

Figure 11 shows a series of ontology management services which we are developing. Each of the services can be used independently or the user can choose to use several of them.

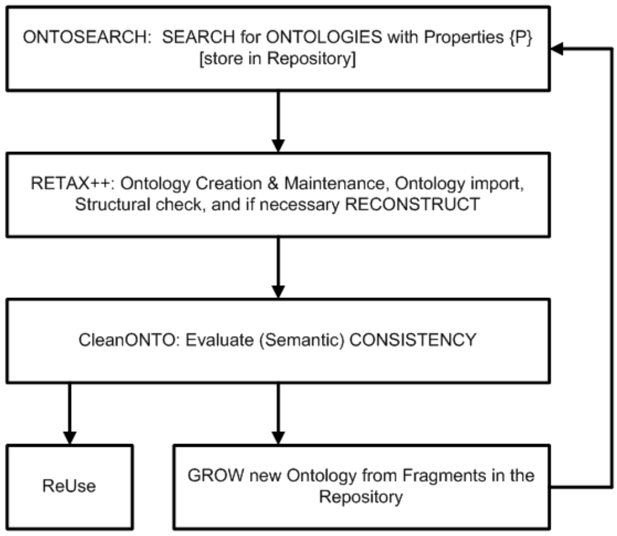

Figure 11: Aberdeen Ontology Management System

For example, ONTOSEARCH is a service which helps a user find ontologies/ontological fragments which contains certain concepts/keywords [16]. ReTAX++ provides a series of checks on the consistency of an ontology and can additionally be used to reconstruct an ontology from a series of fragments [10]. CleanONTO is a further service which detects and removes inconsistencies associated with an ontology's taxonomic structure.

9. ACKNOWLEDGMENTS

This work is funded by the IPAS project (DTI project No. TP/2/IC/6/I/10292). Special thanks to Jeff Pan, Joey Lam and David Fowler for their feedback.

10. REFERENCES

[1] C. Brewster, H. Alani, S. Dasmahapatra, and Y. Wilks. Data driven ontology evaluation. In *Proceedings of International Conference on Language Resources and Evaluation*, Lisbon, Portugal, 2004.

[2] C. Fellbaum. *WordNet: An Electronic Lexical Database*. MIT Press, 1998.

[3] A. Gangemi, N. Guarino, C. Masolo, and A. Oltramari. Sweetening wordnet with dolce. *AI Magazine*, 24(3):13–24, September 2003.

[4] A. Gangemi, N. Guarino, and A. Oltramari. Conceptual analysis of lexical taxonomies: The case of wordnet top-level. In C. Welty and B. Smith, editors, *Proceedings of FOIS 2001*, pages 285–296. ACM press, 2001.

[5] A. Gomez-Perez. Evaluation of ontologies. *International Journal of Intelligent Systems*, 16(3):391–409, 2001.

[6] N. Guarino. The role of identity conditions in ontology design. *Lecture Notes in Computer Science*, 1661:221–227, 1999.

[7] N. Guarino and C. Welty. Evaluating ontological decisions with ontoclean. *Communications of the ACM*, 45(2):61–65, February 2002.

[8] N. Guarino and C. A. Welty. An overview of ontoclean. In S. Staab and R. Studer, editors, *Handbook on Ontologies*, International Handbooks on Information Systems, pages 151–172. Springer, 2004.

[9] A. N. Kaplan. Towards a consistent logical framework for ontological analysis. In C. Welty and B. Smith, editors, *Formal Ontology in Information Systems*. ACM Press, 2001.

[10] S. J. Lam, D. Sleeman, and W. Vasconcelos. Retax+: A cooperative taxonomy revision tool. In *Proceedings of AI-2004 Conference, Cambridge, UK, December 2004*, pages 64–77. Springer, 2004.

[11] G. Miller. Nouns in wordnet: a lexical inheritance system. *International Journal of Lexicography*, 3(4):245–264, 1990.

[12] N. F. Noy and D. L. McGuinness. Ontology development 101: A guide to creating your first ontology. Technical report, Stanford Knowledge Systems Laboratory Technical Report KSL-01-05, 2001.

[13] A. Oltramari, A. Gangemi, N. Guarino, and C. Masolo. Restructuring wordnet's top-level: The ontoclean approach. In *Proceedings of LREC2002 (OntoLex workshop). Las Palmas, Spain*, 2002.

[14] J. Volker, D. Vrandecic, and Y. Sure. Automatic evaluation of ontologies (aeon). In Y. Gil, E. Motta, V. R. Benjamins, and M. A. Musen, editors, *The SemanticWeb - ISWC 2005, 4th International SemanticWeb Conference, ISWC 2005, Galway, Ireland, November 6-10, 2005, Proceedings*, volume 3729, pages 716–731. Springer, 2005.

[15] C. A. Welty and N. Guarino. Supporting ontological analysis of taxonomic relationships. In *Data and Knowledge Engineering*, volume 39, pages 51–74. Elsevier Science B.V., 2001.

[16] Y. Zhang, W. Vasconcelos, and D. Sleeman. Ontosearch: An ontology search engine. In *Proceedings The Twenty-fourth SGAI International Conference on Innovative Techniques and Applications of Artificial Intelligence (AI-2004)*, Cambridge, UK, 2004.