

Ontology Debugging with Truth Maintenance Systems

Hai H. Nguyen, Natasha Alechina and Brian Logan¹

1 Introduction

An ontology is a description of a particular domain in terms of its concepts and relationships. Ontologies can be used by ‘intelligent agents’ to model their environment and communicate with other agents. The quality of ontologies are therefore crucial to the performance of intelligent agents. However, as with any other knowledge base, there is always the possibility that an ontology have some semantic defects. This abstract presents an approach to ontology debugging using ideas borrowed from Truth Maintenance Systems (TMS).

2 Ontology Debugging

This section briefly introduces the key problems in ontology debugging and briefly outlines some of the major work in the field. Firstly, we provide basic definitions of incoherence and inconsistency of a DL-based ontology.

Definition 1. *An ontology is incoherent iff there is at least one unsatisfiable concept in its TBox.*

Definition 2. *An ontology is inconsistent iff there is no model for it.*

Generally speaking, the incoherence problem deals with concept-unsatisfiability within the TBox while an inconsistency problem also involves assertional axioms.² *Ontology debugging* is the process of identifying bugs and producing repair plans for an incoherent or inconsistent ontology. Most work has been done in ontology debugging is for the incoherence problem (i.e., debugging and repair of unsatisfiable concepts) although recently the problem of inconsistency has also been investigated.

Basically, the process of debugging ontology has two parts. The first is to identify which sets of axioms (or parts of axioms) are responsible for an unsatisfiable concept or an inconsistency. The next step is to propose how can these axioms be modified to make the concept satisfiable, or to restore consistency to the Knowledge Base (KB) with respect to some particular criteria.

Two main approaches to pinpointing problematic axioms have been proposed in the literature: glass-box and black-box. Glass-box methods, e.g., [6, 7, 8, 9], use tableau-like rules to pinpoint the problematic axioms (concepts). These approaches obviously depend on a particular DL, as they have to modify the tableau rules to store and retrieve the sources of errors. Black-box methods on the other hand,

e.g., [6, 11, 4], are reasoner-independent, since they only use the reasoner as an external component to diagnose whether an a concept is satisfiable with respect to a particular T-Box (or KB in the case of inconsistency problem). There are also hybrid approaches, e.g., [5], which combine both glass box and black box approaches.

In this paper, the rules which we employ to create the dependency graph are similar to the classical tableau rules, as in some of the glass-box approaches, e.g., [7, 8].

3 Truth Maintenance Systems

Truth Maintenance Systems (TMS), e.g., [3], also known as Belief Revision Systems or Reason Maintenance Systems, play a central role in a style of belief revision called foundational belief revision. A TMS keeps track of dependencies between data to maintain the consistency of a database. A TMS consists of a set datum nodes and the justifications for them. A justification can be considered as a record of an inference, linking a datum node with the datum nodes used to derive it. Using these recorded dependencies, a TMS allows a problem-solver to quickly determine which nodes are “responsible” for belief in a particular datum.

According to [10], a TMS performs three main tasks: 1) given an assertion, find the assertions or assumptions used to derive it; 2) given a set of assumptions, find all the assertions can be derived from them; and 3) delete an assertion and all the consequences which have been derived from it. These tasks are also relevant to the problem of ontology debugging. For example, tracing the sources S_1 and S_2 of the assertions $A(x)$ and $\neg A(x)$, where A is a concept name and x is an individual in the ontology, gives the source of the contradiction (or clash) $S_1 \cup S_2$. Similarly, if one can find a minimal set of assumptions from which the contradictory assertions were derived, the minimal set of axioms which are the cause for the clash can also be identified.³ This set corresponds to a *MUPS* in [9], or a *justification* for concept unsatisfiability defined in [5].

4 Using ATMS for Ontology Debugging

One particular type of TMS is an Assumption-based TMS (ATMS) [1]. In an ATMS, each node is associated with the set of sets of assumptions used to derive it, as well as the datum nodes that constitute its immediate antecedents. These sets of assumptions are termed *environments*, and are always kept minimal and consistent. In this way, backtracking is avoided and multiple solutions can be found at the same time.

³ In the literature of ontology debugging, the idea of tagging an assertion with the axioms used to derive it has also been proposed in [7, 8].

¹ School of Computer Science, University of Nottingham UK, {hhn,nza,bsl}@cs.nott.ac.uk

² Note that the satisfiability checking problem in TBox can be reduced to a consistency checking problem by trying to construct a model for a concept using tableau rules.

In this section, we present an approach to ontology debugging using an ATMS. We focus on the problem of axiom pinpointing for an unsatisfiable concept (i.e., finding a set of axioms responsible for an unsatisfiable concept), and for simplicity, we only consider the unfoldable \mathcal{ALC} TBOX without disjunctions.⁴ We show how the ATMS can be used to detect contradictions and to pinpoint sets of problematic axioms.

As a reasoner can easily detect that a concept is unsatisfiable by a satisfiability check, the key problem is to identify the sources of the unsatisfiability. This is the task of the ATMS. An ATMS node N_{datum} is of the form: $\langle datum, label, justifications \rangle$, where *datum* is an assertion such as $A_i(a)$, *label* is a set of environments (explained below), and *justifications* are the sets of nodes that directly derive N_{datum} . Since there are many ways a datum can be derived, it is possible to have multiple justifications for a particular node. The ATMS distinguishes two special types of datum nodes: assumptions and premises. *Assumptions* are foundational data. Each environment in the label of a (non-assumption) datum node comprises a set of assumptions from which the datum can consistently be derived. *Premises* are similar to assumptions, but are taken to hold universally, and are not explicitly represented in environments. The task of the ATMS is to ensure that each node label is consistent, sound, complete and minimal. As the reasoner informs the ATMS of new datum nodes and justifications, the ATMS *label propagation* algorithms update the labels of previously asserted nodes to remove any subsumed environments (in the case of a normal justification), or any environments which subsume an environment (in the case of a new justification for the distinguished node N_{\perp} which represents contradiction).

The ontology debugging problem can be mapped onto the operations of the ATMS in a straightforward way. Each TBOX axiom is represented by an ATMS assumption. For concreteness, we assume a TBOX $\Gamma = \{ax_1, \dots, ax_n\}$, where each axiom ax_i is of the form: $A_i \sqsubseteq C_i$ and all concept descriptions are in negation normal form (NNF). The assumption that each concept is non-empty, e.g., $A_i(c)$ for some constant c , is represented by an ATMS premise. The reasoner uses standard rules of inference to infer new concept instances from some consistent set of datum nodes (i.e., nodes whose labels do not subsume the label of N_{\perp}). A suitable list of rules that can be used by the reasoner to infer new justifications is shown in Figure 1. The process of creating the dependency graph terminates when no rule can be applied to any node of the graph. At this point, each environment of a node N_{datum} is a minimal set of axioms that can be used to derive *datum*, and the label of N_{\perp} consists of sets of axioms responsible for clashes. In addition, the information given by justifications for nodes can be used to pinpoint the parts of axioms, e.g., concepts causing the clashes.

In conclusion, there is a clear mapping between the functionality provided by the ATMS and the problems of ontology debugging, and we believe that a systematic investigation of the practicality of using an ATMS for ontology debugging is a fruitful direction for future research.

REFERENCES

- [1] Johan de Kleer, ‘An assumption-based TMS’, *Artificial Intelligence*, **28**(2), 127–162, (March 1986).

⁴ With disjunctions, the setting is more complicated. However, disjunctions can be handled in several different ways, e.g., similarly to the extended ATMS described by de Kleer in [2], or by using sub-graphs to deal with the or-branches.

| | |
|---------------------|--|
| \sqsubseteq -rule | If the current node is $N_{A_i(a)} : \langle A_i(a), L, J \rangle$ and $A_i \sqsubseteq C_i \in \Gamma$, then create a new node $N_{C_i(a)} : \langle C_i(a), L', \{(A_i(a))\} \rangle$ where $L' = \{e \cup \{ax_i\} e \in L\}$. If $N_{C_i(a)}$ exists, update its label and justifications. $L_{C_i(a)} = L_{C_i(a)} \cup L'$; $J_{C_i(a)} = J_{C_i(a)} \cup \{(A_i(a))\}$. |
| \sqcap -rule | If the current node is $N_{C_i(a) \sqcap C_j(a)} : \langle C_i(a) \sqcap C_j(a), L, J \rangle$, then create 2 new nodes $N_{C_i(a)} : \langle C_i(a), L, \{(C_i(a) \sqcap C_j(a))\} \rangle$ $N_{C_j(a)} : \langle C_j(a), L, \{(C_i(a) \sqcap C_j(a))\} \rangle$ if they exist, update their labels and justifications. |
| \exists -rule | If the current node is $N_{(\exists s.C)(a)} : \langle (\exists s.C)(a), L, J \rangle$, then create 2 nodes $N_{s(a,b)} : \langle s(a,b), L, \{(\exists s.C)(a)\} \rangle$ $N_{C(b)} : \langle C(b), L, \{(\exists s.C)(a)\} \rangle$ where b is a new individual. |
| \forall -rule | If the current node is $N_{(\forall s.C)(a)} : \langle (\forall s.C)(a), L, J \rangle$ then if there exists $N_{s(a,b)}$, create a node $N_{C(b)} : \langle C(b), L, \{(\forall s.C)(a), s(a,b)\} \rangle$ |
| \perp -rule | If there exist 2 nodes $N_{A_i(a)} : \langle A_i(a), L_1, J_1 \rangle$ $N_{\neg A_i(a)} : \langle \neg A_i(a), L_2, J_2 \rangle$ in the graph, then create a new node $N_{\perp} : \langle \perp, \{e_1 \cup e_2 e_1 \in L_1, e_2 \in L_2\}, \{j_1 \cup j_2 j_1 \in J_1, j_2 \in J_2\} \rangle$. If $N_{\perp} : \langle \perp, L_{\perp}, J_{\perp} \rangle$ exists, update its label and justification: $L_{\perp} = L_{\perp} \cup \{e_1 \cup e_2 e_1 \in L_1, e_2 \in L_2\}$ $J_{\perp} = J_{\perp} \cup \{j_1 \cup j_2 j_1 \in J_1, j_2 \in J_2\}$ |

Figure 1. Rules for creating and updating nodes in the dependency graph

- [2] Johan de Kleer, ‘Extending the ATMS’, *Artificial Intelligence*, **28**(2), 163–196, (1986).
- [3] Jon Doyle, ‘A truth maintenance system’, *Artificial Intelligence*, **12**(3), 231–272, (1979).
- [4] Matthew Horridge, Bijan Parsia, and Ulrike Sattler, ‘Explaining inconsistencies in OWL Ontologies’, in *Proceedings of SUM '09*, pp. 124–137, Berlin, Heidelberg, (2009). Springer-Verlag.
- [5] A. Kalyanpur, B. Parsia, B.C. Grau, and E. Sirin, ‘Justifications for entailments in expressive description logics’, Technical report, University of Maryland, (2006).
- [6] A. Kalyanpur, B. Parsia, E. Sirin, and J. Hendler, ‘Debugging unsatisfiable classes in OWL ontologies’, *Journal of Web Semantics*, **3**(4), 268–293, (2005).
- [7] Joey Sik Chun Lam, Derek H. Sleeman, Jeff Z. Pan, and Wamberto Weber Vasconcelos, ‘A fine-grained approach to resolving unsatisfiable ontologies’, *J. Data Semantics*, **10**, 62–95, (2008).
- [8] Thomas Meyer, Kevin Lee, Richard Booth, and Jeff Z. Pan, ‘Finding maximally satisfiable terminologies for the description logic ALC’, in *Proceedings of AAAI 2006*. AAAI Press, (2006).
- [9] Stefan Schlobach and Ronald Cornet, ‘Non-standard reasoning services for the debugging of description logic terminologies’, in *Proceedings of IJCAI'03*, pp. 355–360. Morgan Kaufmann, (2003).
- [10] Stuart C. Shapiro, ‘Belief revision and truth maintenance systems: An overview and a proposal’, Technical report, SUNY-Buffalo, (1998).
- [11] Hai Wang, Matthew Horridge, Alan Rector, Nick Drummond, and Julian Seidenberg, ‘Debugging OWL-DL ontologies: A heuristic approach’, in *ISWC 2005, LNCS 3729*, pp. 745–757. Springer, (2005).