Towards a Theory of Ontology Repair

— Or Truthfulness Considered Harmful *

Alan Bundy
School of Informatics, University of Edinburgh
A.Bundy@ed.ac.uk

1 Introduction

We start to develop a formal theory to inform mechanisms for *ontology repair*. In ORS[1] we introduced the important but neglected problem of ontology repair (McNeill and Bundy, forthcoming). We can define this problem as follows:

Definition 1 (Ontology Repair:)

$$Repair(O, \phi, \pi) \quad ::= \quad (O \vdash \phi \wedge \ \models \neg\phi) \rightarrow (\pi(O) \nvdash \pi(\phi) \vee \ \models \pi(\phi)) \tag{1}$$

This definition is then to be read as follows: π is a repair of ontology O wrt formula ϕ iff when ϕ is a theorem of O but is false in the real world, then either $\pi(\phi)$ is not a theorem of $\pi(O)$ or $\pi(\phi)$ is true in the real world[2].

In ORS, the role of ϕ, in the above definition, was played by theorems essentially of the form $O \vdash Holds(goal, sit)$, where sit is a situation in which $goal$ is true according to ontology O. Situation sit can also be read as a *plan* to achieve $goal$. Any subsequent failure of this plan to achieve $goal$ can be represented as $\not\models Holds(goal, sit)$.

For current purposes an *ontology* consists of two parts: a sorted, first-order logical signature, defining some functions and predicates with their types; and a sorted, first-order theory, listing some axioms written in that signature. However, the ideas in this paper are intended to be applicable to a wide range of different logical theories.

We contrast the ontology repair problem with the *ontology matching* problem, in which a theorem-preserving relationship is to be found between *two* different ontologies, without adapting either of them. Ontology repair, on the other hand, adapts a *single* ontology by reference to its intended semantics.

2 Ontology Repair Operations

In the ORS system, we considered 9 ontology repair operations. These were as follows.

Precondition Abstraction and Refinement: A precondition of a rule is either dropped (abstraction) or added (refinement).

Propositional Abstraction and Refinement: An argument of a predicate is either dropped (abstraction) or added (refinement).

Predicate Abstraction and Refinement: Two or more predicates are merged (abstraction) or a predicate is split into two or more (refinement).

Domain Abstraction and Refinement: The terms of a domain are merged into equivalence classes (abstraction) or terms are split up into a class of objects of the same type (refinement).

Argument Permutation: The arguments of a predicate are permuted.

The four kinds of abstraction are all those identified in the extensive survey of applications of abstraction in automated reasoning in (Giunchiglia and Walsh, 1992). The four kinds of refinement were constructed by inverting the four abstractions. Argument permutation was the only other form of repair for which we detected a need during the ORS work.

*Thanks to Fiona McNeill and Andrew Priddle-Higson for feedback on earlier drafts of this paper and to other members of the Mathematical Reasoning Group for feedback on a talk about this work. The work is part of the OpenKnowledge Project, funded by European Commission 6th Framework (IST).

[1] Ontology Repair System

[2] In an abuse of notation, which we hope aids readability, we have overloaded π as a function of both ontologies and sentences.

3 Why Truthfulness is Considered Harmful

A truthful[3] repair operation is one in which $O \vdash \phi \implies \pi(O) \vdash \pi(\phi)$, i.e., in which a theorem in the original ontology is mapped to a theorem in the repaired one. For all truthful repair operations the definition of ontology repair collapses to:

$$Repair(O, \phi, \pi) \quad ::= \quad (O \vdash \phi \wedge \ \models \neg\phi) \implies \models \pi(\phi)$$

That is, since the false conjecture ϕ cannot be rendered unprovable by a truthful repair, we have no choice but to try to make it true.

Unfortunately, it is in the nature of many repair operations that the semantics of ϕ is essentially the same as that of $\pi(\phi)$, so that $\not\models \phi \implies \not\models \pi(\phi)$. This means that there is no hope of verifying the repair operation. So, truthfulness is harmful in this context, and should be avoided if possible.

4 Propositional Abstraction and Refinement

Space precludes a discussion of all nine ontology repair operations, so we confine ourselves to two of them: propositional abstraction and refinement, but focus on some of the generic issues that arise for all the operations. The straightforward way to implement these two operations is a simple syntactic deletion or addition of an argument to each occurrence of the repaired predicate in the signature and theory of the ontology. However, this straightforward approach is problematic in several ways.

- Propositional abstraction is truthful (see (Giunchiglia and Walsh, 1992) for a proof).
- Propositional abstraction can turn consistent ontologies into inconsistent ones, e.g. if both $p(a) \in Ax(O)$[4] and $\neg p(b) \in Ax(O)$ then $\pi(O) \vdash p \wedge \neg p$.
- In propositional refinement the value to be given to the additional argument is not always clear.

Similar problems occur with simple syntactic definitions of the other repair operations.

We, therefore, propose alternative semantic definitions that do not suffer from such problems. For instance, for propositional abstraction and refinement, respectively, we propose the definitions.

$$\begin{aligned} Ax(\pi_a(O)) \quad &::= \quad \{\phi[\overline{p(\bar{t})}] \mid \exists \overline{t_0}.\ \phi[\overline{p(t_0, \bar{t})}] \in Ax(O) \wedge O \not\vdash \neg\phi[\overline{\exists x. p(x, \bar{t})}]\} \\ Ax(\pi_r(O)) \quad &::= \quad \{\phi[\overline{\exists x. p(x, \bar{t})}] \mid \phi[\overline{p(\bar{t})}] \in Ax(O)\} \end{aligned}$$

where $\bar{t}$ is a vector of values. Note now that:

$$\neg p \notin Ax(\pi_a(O)) \iff \neg\exists t_0.\ \neg p(t_0) \in Ax(O) \vee O \vdash \neg\neg\exists x.\ p(x) \iff O \vdash \exists x.\ p(x)$$

So, $\neg p \notin Ax(\pi_a(O))$ even though $\neg p(b) \in Ax(O)$ because $p(a) \in Ax(O)$. However, $p \in Ax(\pi_a(O))$ since $O \not\vdash \neg\exists x.\ p(x)$.

5 Conclusion

We have given a formal definition of the problem of ontology repair and are working towards formal definitions of some ontology repair operations. We note that truthfulness in an ontology repair operation is an undesirable property as it makes it difficult, if not impossible, then to use it for ontology repair. To avoid truthfulness, and related problems of inconsistency and definitional vagueness, in the straightforward syntactic definitions of these repair operations, we are experimenting with semantic alternatives. We plan to prove theoretic properties of these semantically defined repair operations and to implement them and evaluate them experimentally.

References

F. Giunchiglia and T. Walsh. A theory of abstraction. *Artificial Intelligence*, 56(2–3):323–390, 1992. Also available as DAI Research Paper No 516, Dept. of Artificial Intelligence, Edinburgh.

F. McNeill and A. Bundy. Dynamic, automatic, first-order ontology repair by diagnosis of failed plan execution. *IJSWIS*, forthcoming. Special issue on ontology matching.

[3]This was Giunchiglia and Walsh's original notation. They later changed it to TI (for Theorem Increasing), but "TI considered harmful" does not have the same cachet.

[4]Where $Ax(O)$ is the set of axioms of ontology O.