Edinburgh Research Explorer

Unite: A New Plan for Automated Ontology Evolution in Physics

Citation for published version:
Bundy, A 2009, Unite: A New Plan for Automated Ontology Evolution in Physics. in *Notes of the IJCAI-09 Workshop ARCOE-09.*

Link:
Link to publication record in Edinburgh Research Explorer

Document Version:
Early version, also known as pre-print

Published In:
Notes of the IJCAI-09 Workshop ARCOE-09

Unite: A New Plan for Automated Ontology Evolution in Physics[*]

Alan Bundy

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

Abstract

We are developing a novel technique for ontology evolution, which we call *ontology repair plans*. Our development case studies are drawn from instances in the history of physics where experimental observation contradicted current physical theories, which then had to be evolved. In particular, it was often necessary to evolve the *representation language* of these theories, and not just the physical laws. To date, we have implemented one ontology repair plan for splitting a function into three and another for adding an additional argument to a function. In this paper, we describe a new ontology repair plan *Unite*, for equating two previously distinct functions.

1 Ontology Repair Plans

We are developing a series of *ontology repair plans* which operate simultaneously on a small set of modular higher-order[1] ontologies, e.g., one representing an initial theory of physics, another representing a particular experimental set-up [Bundy and Chan, 2008]. Each repair plan has a trigger formula and some actions: when the trigger is matched, the actions are performed. The actions modify both the signatures[2] and the axioms of the old ontologies to produce new ones. The repair plans have been implemented in the GALILEO system (Guided Analysis of Logical Inconsistencies Leads to Evolved Ontologies) using λProlog [Miller and Nadathur, 1988] as our implementation language, because it provides a polymorphic, higher-order logic programming language.

We have so far developed two repair plans, which we call *Where's my stuff?* (WMS) and *Inconstancy*. These roughly correspond to the operations of splitting a function and adding an argument, respectively. We have found multiple examples of these repairs across the history of physics.

The WMS repair plan aims at resolving contradictions arising when the predicted value returned by a function does not match the observed value. This is modelled by having two ontologies, corresponding to the prediction and the observation, with different values for this function. To break the inconsistency, the conflicted function is split into three new functions: *visible*, *invisible* and *total*. The conflicted function becomes the total function in the predictive theory and the visible function in the observation theory[3]. The invisible function is defined as the difference between them, and this new definition is added to the predictive theory. The intuition behind this repair is that the discrepancy arose because the function was not being applied to the same *stuff* in the predictive and the observational ontologies — the invisible stuff was not observed.

WMS has been successfully applied to conflicts between predictions of and observations of the following functions: the temperature of freezing water; the energy of a bouncing ball; the graphs relating orbital velocity of stars to distance from the galactic centre in spiral galaxies; and the precession of the perihelion of Mercury. In these examples the role of the invisible stuff is played by: the latent heat of fusion, elastic potential energy, dark matter and an additional planet, respectively.

The Inconstancy repair plan is triggered when there is a conflict between the predicted independence and the observed dependence of a function on some parameter, i.e., the observed value of a function unexpectedly varies when it is predicted to remain constant. This generally requires several observational ontologies, each with different observed values of the function, as opposed to the one observational theory in the WMS plan. To effect the repair, the parameter causing the unexpected variation is first identified and a new definition for the conflicted function is created that includes this new parameter. The nature of the dependence is induced from the observations using curve-fitting techniques.

Inconstancy has been successfully applied to the following conflicts between predictions and various observations: the ratio of pressure and volume of a gas; replacement of Aristotle's concept of instantaneous light travel with a finite (but fast) light speed; and again the graphs relating orbital velocity of stars to distance from the galactic centre in spiral galaxies. The unexpected parameter of the function is: the temperature

---

[*]I'm grateful to Michael Chan and Jos Lehmann for comments on an earlier draft, and to both Michael and Alan Smaill for programming assistance.

[1]The physics domain requires higher-order logic: both at the object-level, to describe things like planetary orbits and calculus, and at the meta-level, to describe the ontology repair operations.

[2]A signature describes the representation language of an ontology, e.g., its functions and their types.

[3]There are situations in which these roles are inverted.

of the gas; the distance between the source and target of a flash of light; and the acceleration between the stars, respectively. The first of these repairs generalises Boyle's Law to the Ideal Gas Law, the second generalises a moment of light travel to an interval and the third generalises the Gravitational Constant to Milgrom's MOND (MOdified Newtonian Dynamics). Interestingly, WMS and Inconstancy produce the two main rival ontologies on the spiral galaxy anomaly, namely dark matter and MOND. Since this is still an active controversy, its unfolding will help us develop mechanisms to choose between rival theory repairs.

The merging of functions and the dropping of arguments are identified as two common forms of *abstraction* in [Giunchiglia and Walsh, 1992]. The abstractions and their inverses, the *refinements* of *splitting* functions *adding* arguments, were used as the basis for ontology evolution in ORS [McNeill and Bundy, 2007]. The GALILEO work addresses one of the key outstanding issues in ORS: the essential ambiguity of refinement. When a function is split it is necessary to decide which occurrences of the old function map to which new function. Similarly, it is necessary to decide the value of each occurrence of a new argument. GALILEO addresses this problem by making uniform mappings within each ontology.

Function merging and argument dropping are examples of *signature* evolution, i.e., a change in the underlying language of an ontology. This is complementary to *belief revision*, which manages axiom evolution within a fixed signature. Both are needed. Below, we explore an ontology repair plan that creates a new axiom.

2 The Unite Ontology Repair Plan

We now explore the *converse* of the WMS ontology repair plan, which we will call *Unite*. The idea of the Unite plan is to take two different functions and equate them. The classic example of the need for this repair is in the identification of the Evening Star and the Morning Star as two manifestations of Venus. In this case the functions are all nullary, i.e., constants, but we will also give a non-nullary example.

2.1 Defining Properties

If two different terms refer to the same thing then they should yield the same value for each of their properties. Unfortunately, this is impractical as a trigger for the repair plan, because, in practice, the values of *all* properties of a thing are unlikely to be known — indeed, some of the properties themselves may not yet be known. Fortunately, there are properties which *alone* are sufficient. I will call these *defining properties*. For instance, for physical objects, such as heavenly bodies, their *orbit* is a defining property, which we can summarise as saying that two different objects cannot be at the same place at the same time. Note that I mean 4D orbits, not mere 3D ones, i.e., it is not just that the orbits occupy the same path in 3D space, but that each moment of time defines the same position in both orbits.

We will use $DefProp(dp, \tau)$ to represent that property dp is a defining property of objects of type τ, i.e.,

$$DefProp(dp, \tau) \quad ::= \quad \forall x, y{:}\tau.\ dp(x) = dp(y) \rightarrow x = y$$

We are now ready to formally define the Unite repair plan.

2.2 The Plan Formalism

Trigger: If $stuff_1$ and $stuff_2$ are functions of the same type, not already known to be equal, that are observed to take the same value for some defining property, dp, for functions of type τ then the following formula will be triggered.

$$
\begin{aligned}
O_t &\not\vdash \quad stuff_1 = stuff_2, \\
O_t &\vdash \quad stuff_1 : \tau \wedge stuff_2 : \tau \wedge DefProp(dp, \tau), \\
O_s &\vdash \quad dp(stuff_1) = dp(stuff_2). \qquad\qquad (1)
\end{aligned}
$$

O_t is the theoretical ontology, O_s is the local ontology describing a particular set of experiments.

Create New Ontologies: The repair is to add an equality between $stuff_1$ and $stuff_2$ as a new axiom to O_t. O_s is unchanged.

$$Ax(\nu(O_t)) \quad ::= \quad \{stuff_1 = stuff_2\} \cup Ax(O_t) \ (2)$$

where $Ax(O)$ is the set of axioms of ontology O and $\nu(O)$ is the ontology resulting from repairing O.

2.3 The Plan Implementation

Unite has been implemented in the GALILEO system. This implementation consists of λProlog code for:

- an additional clause in the generic `Repair` function, that takes two ontologies and outputs their repaired axioms;
- a function `unite_trigger` that checks that the trigger formula holds; and
- a function `changeUnite` that adds the additional axiom to the theoretical ontology.

Note that the physical formulae are represented using a *deep embedding*, in which `applic` applies a function to an argument and `lambda` abstracts a variable in a function. `turnstile O T` represents $O \vdash T$. `equal` is defined as a unary function on a list of equal objects. In λProlog, variables start with an upper case letter; constants with lower case[4].

```
% Unite repair
repair Ot Os NAt As :-
 unite_trigger Ot Os S1 S2,
 changeUnite Ot S1 S2 NAt,
 axioms Os As.

% Unite trigger:
unite_trigger Ot Os S1 S2 :-
 not turnstile Ot (applic equal (S1::S2::nil)),
 turnstile Ot (applic (applic isa S1) T),
 turnstile Ot (applic (applic isa S2) T),
 turnstile Ot (applic (applic defprop DP) T),
 turnstile Os (applic equal ((applic DP S1)
  ::(applic DP S2)::nil )).

% Unite repair:
changeUnite O S1 S2 ((applic equal
    (S1::S2::nil))::A) :-
 axioms O A.
```

---

[4]Unfortunately, the opposite of the normal mathematical conventions used elsewhere in this paper.

3 Case Studies

We now describe two case studies that were used as the development set for Unite, and have been successfully implemented and evaluated.

3.1 Example: The Morning and Evening Stars

Because Venus is closer to the Sun than the Earth, it becomes visible either just before dawn or just after sunset, when it is the brightest heavenly object after the Moon. These two kinds of appearance were not originally identified as coming from the same object. Both the Ancient Egyptians and the Ancient Greeks thought there were two objects. The pre-dawn appearances were identified as the Morning Star and the post-sunset ones as the Evening Star. It was only with the quantification of astronomy that the orbits of these 'two' 'stars' were calculated and seen to be the same (up to experimental error).

We can use the Unite repair plan to emulate this episode as follows. The trigger formulae are:

$$
\begin{aligned}
O_t &\nvdash \; MS = ES \\
O_t &\vdash \; MS : obj \wedge ES : obj \wedge DefProp(orbit, obj) \\
O_s &\vdash \; orbit(MS) = orbit(ES).
\end{aligned}
$$

where MS and ES are constants standing for "Morning Star" and "Evening Star", respectively. These formulae match with (1). The repair is then:

$$
Ax(\nu(O_t)) \;\; ::= \{MS = ES\} \cup Ax(O_t)
$$

from (2), as required. When the objects being equated are constants, as in this example, we might want to go further and replace both old constants with a new one, e.g. $Venus$, but this is beyond the scope of this paper.

3.2 Example: The Earth as a Sphere

Pythagoras was one of the first astronomers to realise that the Earth was a sphere. He gathered evidence to support this theory from various sources, but in this paper we will consider only his observations of lunar eclipses. He noticed that the edge of the shadow that the Earth cast on the Moon was always circular. He reasoned that the only 3D shape that always casts circular shadows is a sphere.

This reasoning is also an example of the Unite repair plan. This time the terms being equated are compound: $Shape(Earth)$ and $Shape(Ball)$, where $Ball$ is some imaginary spherical object in the same orbit as the Earth. The defining property is $\lambda v, t. \; project(v, Sun, Moon, t)$: the projection of a volume v from the Sun onto the Moon. The idea is that if two 3D objects always have the same 2D projections then they have the same shape, i.e., $DefProp(\lambda v, t. \; project(v, Sun, Moon, t), vol)$.

Note that multiple, independent projections are required. A cylinder also projects as a circle along its axis, but most of its projections are not circular, so one projection is not enough. The abstraction over time supplies these. Pythagoras could not, of course, observe all possible lunar eclipses, so an element of induction is required in the observational ontology O_s. Note also that the 'observed' projections of $Ball$ are a thought experiment.

The application of Unite proceeds as follows. The trigger formulae are:

$$
\begin{aligned}
O_t &\nvdash \; Shape(Earth) = Shape(Ball) \\
O_t &\vdash \; Shape(Earth) : vol \wedge Shape(Ball) : vol \\
&\qquad \wedge DefProp(\lambda v, t. \; project(v, Sun, Moon, t), vol), \\
O_s &\vdash \; \lambda t. \; project(Shape(Earth), Sun, Moon, t) \\
&\qquad = \lambda. \; project(Shape(Ball), Sun, Moon, t)
\end{aligned}
$$

where

$$
\begin{aligned}
&\lambda t. \; project(Shape(Earth), Sun, Moon, t) \\
&\quad = \lambda t. \; project(Shape(Ball), Sun, Moon, t)
\end{aligned}
$$

is Pythagoras' generalisation from his observations of lunar eclipses and his thought experiments. The repair to O_t is to add the new axiom $Shape(Earth) = Shape(Ball)$.

4 Conclusion

In this paper we have described Unite, a converse repair plan to WMS, which equates two different terms. Its central notion is of a defining property, i.e., a property whose value is unique for things of a specific type. We have given a formal definition of the trigger formulae and of the repair operations. This repair plan has been implemented in the GALILEO ontology evolution system. We have illustrated the operation of the Unite plan with two examples:

- the identification of the Morning Star and the Evening Star, where the defining property is their orbit; and

- the identification of the shape of the Earth as that of a sphere, where the defining property is the projection of a 3D volume onto a 2D surface.

These two examples have been successfully evaluated in GALILEO. In future work, we plan to evaluate Unite on a test set of case studies from the history of physics. Since the repair is to add a new law to a physical theory, it has the potential to emulate many major advances in physical ontology evolution.

References

[Bundy and Chan, 2008] A. Bundy and M. Chan. Towards ontology evolution in physics. In W. Hodges, editor, *Procs. Wollic 2008*. LNCS, Springer-Verlag, July 2008.

[Giunchiglia and Walsh, 1992] F. Giunchiglia and T. Walsh. A theory of abstraction. *Artificial Intelligence*, 56(2–3):323–390, 1992.

[McNeill and Bundy, 2007] F. McNeill and A. Bundy. Dynamic, automatic, first-order ontology repair by diagnosis of failed plan execution. *IJSWIS*, 3(3):1–35, 2007. Special issue on ontology matching.

[Miller and Nadathur, 1988] D. Miller and G. Nadathur. An overview of λProlog. In R. Bowen, editor, *Proceedings of the Fifth International Logic Programming Conference/ Fifth Symposium on Logic Programming*. MIT Press, 1988.