



THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Evolution of Inconsistent Ontologies in Physics

Citation for published version:

Lehmann, J, Bundy, A & Chan, M 2011, Evolution of Inconsistent Ontologies in Physics. in *Proceedings of the Symposium on Inconsistency Robustness 2011*.

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Early version, also known as pre-print

Published In:

Proceedings of the Symposium on Inconsistency Robustness 2011

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.



Evolution of Inconsistent Ontologies in Physics^{*†}

Jos Lehmann, Alan Bundy, Michael Chan

School of Informatics, University of Edinburgh, Edinburgh, UK

{jlehmann, bundy, mchan}@inf.ed.ac.uk

Abstract

Inconsistency robustness in autonomous software can be seen as a problem of automated reasoning about ontology evolution. Formal ontologies specify the knowledge that software systems use when reasoning about the entities in their domain. Such knowledge is bound to evolve in the face of new information. Robust software should therefore be able to maintain the consistency between its own ontologies and any incoming information that contradicts them. This can be achieved either by isolating the inconsistency or by evolving the ontologies.

We propose a higher-order logical approach to ontology evolution and apply it to examples in physics, as advances in this field are naturally modelled as cases of ontology evolution. GALILEO, a system based on this approach, is being implemented and tested. Its basic mechanisms for evolution are *ontology repair plans*. These operate on ontologies formalised and implemented as contexts, which are logical theories that use their own local concepts to describe the domain, thus preventing potential contradictions with other theories to arise. When, though, ontologies are mapped or aligned, they share axioms. This may allow the proof of contradictory facts that affect the robustness of the system. At this stage, the application of an ontology repair plan may resolve the inconsistency, as each plan compiles together a pattern for diagnosis of conflicts between ontologies and transformation rules for effecting a repair. The repair can combine the retraction of axioms, the change of beliefs as well as the deeper modification of the language in which the ontology is represented.

1 Introduction

Artificial intelligence and, more generally, computer science are presently faced with the challenge of how autonomous software can achieve inconsistency robustness by

^{*}The research reported was supported by EPSRC grant EP/E005713/1 and an ORSAS award.

[†]Copyright of this paper belongs to the authors, who hereby grant permission to Stanford University to include it in the conference material and proceedings of Inconsistency Robustness 2011 and to place it on relevant websites.

manipulating its own knowledge. Such knowledge is typically represented in an ontology that conceptualises the entities of the software's application domain and allows the software to reason about such entities at a higher level of abstraction than simply the level of data or information. Just like any abstract model, ontologies are limited representations of the world, which is dynamic and inherently complex. If autonomous systems are to feature any kind of robustness with respect to the dynamics and complexity of changing environments and goals, of communication acts and of new information, they must be able to autonomously update their own ontologies.

The literature on the subject of updating an ontology in the face of new information often uses the phrase *ontology evolution* and usually concentrates on how Description Logic (DL) axiomatic theories for Semantic Web applications need to retract axioms or modify entailments in order to maintain their coherence and consistency. Section 2 discusses work on ontology evolution and inconsistency robustness that is related to ours.

Based on this discussion, we present our approach to formalizing and automating ontology evolution in Higher Order Logic (HOL), which underlies the GALILEO¹ system. In GALILEO, the basic mechanisms for evolution are called *ontology repair plans* (ORPs). Each ORP compiles together a pattern for diagnosis of conflicts between ontologies and transformation rules for effecting a repair. For both development and testing, we rely on examples from physics, as advances in this field may naturally be modelled as cases of ontology evolution and they are usually well documented. Physicists revise predictive theories when confronted with conflicting experimental evidence. Therefore, the ORPs typically assume there is one ontology representing a predictive theory and a second ontology representing an experimental or observational set-up for that theory. When the experimental ontology contains a theorem that contradicts one of the theoretical ontology, an ORP is triggered and amends the two ontologies. The development methodology of GALILEO revolves around the selection of initial ontologies by the collection, analysis, formalisation, implementation and testing of appropriate case studies in the history of physics and of the ORPs inspired by them.

The initial ontologies are formalized and implemented as contexts, i.e., as logical theories which are isolated from other theories and which use their own local concepts to describe the entities in the domain. In such a setting, ontologies may *implicitly* contradict one another without though producing an *explicit* logical contradiction at the global level. When, though, two or more such ontologies are *bridged*, i.e., mapped or aligned through a third one which merges them and resolves their differences, they will be able to share axioms. This will allow the proof of contradictory facts thus affecting the robustness of the system. In this type of situation we say that the ontologies are *locally consistent* but *globally inconsistent*. And it is when the global inconsistency has become explicit that the application of an ORP may resolve the contradiction and re-establish consistency.

Shortly, we investigate in Higher Order Logic the problem of ontology evolution from the perspective of automating the mechanisms to repair locally consistent but globally inconsistent ontologies.

The empirical part of our methodology supports the long-term objectives of the

¹Guided Analysis of Logical Inconsistencies Leads to Evolved Ontologies.

definition of a theory of ontology evolution, as well as of an evolution calculus. The overall aim is to demonstrate that automatic ontology evolution via ORPs is computationally feasible and can account for the kinds of ontology evolution that are observed in human problem solving in the physics domain, and possibly in similar ones. We would like to achieve desirable properties, e.g., coverage, efficiency, maintainability, high quality of the repairs. So far we have been able to evaluate the generality of a higher-order logical approach to ontology evolution and the meaningfulness of the possible evolutions proposed by GALILEO.

Section 2 discusses related work. Section 3 describes the adopted methodology. Sections 4, 5, 6 present three ORPs and applications for each them. These three sections are quite dense in formal content, they describe in detail the formal apparatus of our proposal as well as a number of higher-order logical models of physics case studies to which our ORPs apply. Section 7 briefly illustrates the initial implementation of GALILEO. Section 8 draws some conclusions and discusses initial elements of evaluation.

2 Discussion of related work

The literature on the subject of updating an ontology in the face of new information often uses the phrase *ontology evolution* and usually concentrates on how Description Logic (DL) ontologies for Semantic Web applications need to evolve, either to maintain their own coherence and consistency (ontology debugging) or to establish a relationship with other ontologies (ontology alignment). Debugging yields notions like incoherence and inconsistency diagnosis and repair [Haase *et al*, 2005, Kalyanpur *et al*, 2006b, Ji *et al*, 2009, Lam *et al*, 2008], belief revision [Flouris, 2006], conservative extensions [Ghilardi *et al*, 2006]; alignment yields notions like matching [Doan *et al*, 2004, Giunchiglia & Shvaiko, 2004], mapping and contextualisation [Bouquet *et al*, 2004].

A standard example used in the DL literature about incoherence and inconsistency diagnosis and repair is shown in Figure 1. Note the following DL conventions: $A \sqsubseteq B$ is the DL notation for the first-order logic formula $\forall x.A(x) \rightarrow B(x)$; $i : A$ is the DL notation for i is an instance of A ; a DL ontology consists of a TBox, \mathcal{T} , in which concepts are specified using the signature elements of $sig(\mathcal{T})$, and of an ABox, \mathcal{A} where assertions are made about the individuals using the terminology. The axiomatic theory of members of a university formed by the ordered couple $(\mathcal{T}_{uni}, \mathcal{A}_{uni})$ is a knowledge base. In DL terms [Flouris *et al*, 2006], $(\mathcal{T}_{uni}, \mathcal{A}_{uni})$ is *inconsistent*, i.e., it has no models, because axioms (7, 5) allow to conclude that *bruce* is a *Student* while axioms (7, 6, 4) allow to conclude that *bruce* is not a *Student*. The source of the inconsistency lies in \mathcal{T}_{uni} , which is *incoherent*, i.e., it contains the *unsatisfiable* concept *PhDStudent*. A concept is unsatisfiable if it is mapped to the empty set in all models of the ontology.

In order to diagnose the unsatisfiable concept and repair the incoherence and the inconsistency, DL approaches such as [Haase *et al*, 2005] try to identify the axiom that, if removed, would allow to re-establish coherence and consistency. This is done by algorithms that compute so-called *Minimal Unsatisfiability-Preserving Subsets* of the ontology (MUPS) as well as its *Minimal Inconsistent Subsets* (MIS). For instance, start-

$$\begin{aligned}
\text{sig}(\mathcal{T}_{uni}) &::= \{ \\
&\dots, \text{Person}, \text{Student}, \text{Employee}, \dots \} & (1) \\
\mathcal{T}_{uni} &::= \{ \\
&\text{Student} \sqsubseteq \text{Person}, & (2) \\
&\text{Employee} \sqsubseteq \text{Person}, & (3) \\
&\text{Employee} \sqsubseteq \neg \text{Student}, & (4) \\
&\text{PhDStudent} \sqsubseteq \text{Student}, & (5) \\
&\text{PhDStudent} \sqsubseteq \text{Employee} \} & (6) \\
\mathcal{A}_{uni} &::= \{ \\
&\text{bruce} : \text{PhDStudent} \} & (7)
\end{aligned}$$

Figure 1: Example of an inconsistent DL ontology of members of a university. Axioms (7, 5) allow to conclude that *bruce* is a *Student* while axioms (7, 6, 4) allow to conclude that *bruce* is not a *Student*. The concept *PhDStudent* is unsatisfiable.

ing from a subset of the ontology containing one axiom, the algorithm adds to the subset, one at the time, the axioms that are *connected* to it, i.e., that share type symbols with it (e.g. axioms (4, 5)), until the subset is unsatisfiable. In Figure 1 the only MUPS of \mathcal{T}_{uni} contains axioms (4, 5, 6). A judgement call is therefore required to decide which axiom to remove. This decision procedure may dramatically increase in complexity in real-world scenarios, where MUPS are often more than one.

Even in cases where only one MUPS is returned, the axiomatic structure of the ontology may be entangled and make axiom-removal simply infeasible. Consider for instance the version of $(\mathcal{T}_{uni}, \mathcal{A}_{uni})$ given below (where $/$ is set subtraction) in which axioms (4, 5) are substituted by one single axiom (8):

$$\begin{aligned}
\mathcal{T}_{uni}^* &::= \mathcal{T}_{uni} / \{ \text{Employee} \sqsubseteq \text{Person}, \text{Employee} \sqsubseteq \neg \text{Student} \} \cup \{ \\
&\text{Employee} \sqsubseteq \text{Person} \sqcap \neg \text{Student} \} & (8)
\end{aligned}$$

Removing axiom (8) would affect the entailment structure of \mathcal{T}_{uni}^* , beyond what is needed and make the repair harmful, as an instance of *Employee* would no more qualify as an instance of *Person*. This has motivated approaches such as [Lam *et al*, 2008, Kalyanpur *et al*, 2006a, Schlobach & Cornet, 2003] that try to refine both the diagnostic and the repair operations, in order to allow for more surgical removals or even axiom rewriting. In such approaches the diagnosis usually applies minimal unsatisfiability to abstract versions of the axioms at hand. For instance, axiom (8) would have to be broken up along its conjunction, i.e., reduced to (4, 5), so that a separate unsatisfiability test can be run on each conjunct. A problematic aspect of this approach to diagnosis is that different types of formula and connective need different types of abstraction rules and, of course, the complexity of the rules increases with the complexity of the axioms. Abstraction may also become a problem when repairing, because the user needs to be

$$\mathcal{T}_{uni}^{**} ::= \mathcal{T}_{uni} / \{Employee \sqsubseteq \neg Student\}$$

$$\mathcal{A}_{uni}^{**} ::= \mathcal{A}_{uni}$$

$$\begin{aligned} sig(\mathcal{T}_{com}) ::= \{ \\ \dots, Person, Trainee, Staff, \dots \} \end{aligned} \tag{9}$$

$$\mathcal{T}_{com} ::= \{ \\ Trainee \sqsubseteq Person, \tag{10}$$

$$Staff \sqsubseteq Person, \tag{11}$$

$$Staff \sqsubseteq \neg Trainee \} \tag{12}$$

$$\begin{aligned} Map(\mathcal{T}_{uni}^{**}, \mathcal{T}_{com}) ::= \{ \\ \langle 1, Student, Trainee \rangle, \tag{13} \end{aligned}$$

$$\langle 2, Employee, Staff \rangle \} \tag{14}$$

Figure 2: A locally consistent and globally inconsistent network of DL ontologies. The network formed by $(\mathcal{T}_{uni}^{**}, \mathcal{A}_{uni}^{**})$ and $(\mathcal{T}_{com}, \mathcal{A}_{com})$ through the mapping $Map(\mathcal{T}_{uni}^{**}, \mathcal{T}_{com})$ is incoherent and inconsistent, as it allows to prove that *bruce* is and is not a *Student*. Coherence and consistency are re-established by removing lines (13) or (14) from the mapping.

able to relate back the abstracted axioms to the original ones in order to proceed to their removal, and this is not always a straightforward task.

An additional type of incoherence and inconsistency diagnosis and repair is presented in [Ji *et al*, 2009]. This applies a MUPS- and MIS-based, axiom-removal approach to the case of so-called ontology networks, in particular to the diagnosis and repair of networked ontologies that are locally consistent and globally inconsistent. For instance, in Figure 2, $(\mathcal{T}_{uni}^{**}, \mathcal{A}_{uni}^{**})$ is a coherent and consistent version of $(\mathcal{T}_{uni}, \mathcal{A}_{uni})$. $(\mathcal{T}_{uni}^{**}, \mathcal{A}_{uni}^{**})$ is mapped onto another coherent and consistent ontology $(\mathcal{T}_{com}, \mathcal{A}_{com})$ for company members. The network formed by $(\mathcal{T}_{uni}^{**}, \mathcal{A}_{uni}^{**})$ and $(\mathcal{T}_{com}, \mathcal{A}_{com})$ through the mapping $Map(\mathcal{T}_{uni}^{**}, \mathcal{T}_{com})$ (where $\langle id, t_i, t_j \rangle$ means that term t_i in \mathcal{T}_{uni}^{**} is mapped onto term t_j in ontology \mathcal{T}_{com}) is incoherent and inconsistent, as it allows to prove again that *bruce* both is and is not a *Student*. Coherence and consistency are re-established by removing lines (13) or (14) from the mapping.

In general, we think that the types of approaches described above have proven insufficient to address the study of the *automation* of ontology evolution, because of a number of circumstances. Firstly, most proposed approaches depend on users' instructions, which does not include ontology evolution performed *at runtime*, for instance, by agents in heterogeneous environments. Secondly, automated ontology repair systems, e.g., Swoop [Kalyanpur *et al*, 2006b], or even those that take heterogeneity into account like RaDON [Ji *et al*, 2009], focus mostly on retracting axioms and modifying entailments. Systems like [Lam *et al*, 2008] that go beyond axiom retraction and that

try to enable refined repair operations seem to be doing so on a limited logical basis and they still have to rely on users' judgements in order to evaluate the helpfulness and harmfulness of repairs. Furthermore, such systems do not support deeper syntactic manipulations, e.g., changes of the signature of the language. Thirdly, the focus on DL ontologies does not allow for a sufficiently general analysis and resolution of ontological faults. The limited expressivity of first-order logic, let alone fragments of it such as DL, constitutes a limit on the modelling of ontology evolution. Without the means to quantify over and to reason about the predicates, it is virtually impossible to formalise and automate *sufficiently generic* ontology evolution patterns.

In order to deal with the first two limitations found in the DL-based literature, we investigate the problem of ontology evolution from the perspective of *automating* the mechanisms to repair *locally consistent* but *globally inconsistent* ontologies. As mentioned in Section 1, we assume that the robustness of a system is achieved by having information about the same entities distributed in the system across different ontologies or contexts. When two such contexts are bridged, i.e., mapped through a third one which resolves their differences, they will be able to communicate and prove facts in terms of each other. Such a multiple-ontologies approach has inherent advantages for automating ontology evolution. As a matter of fact, a single inconsistent ontology allows to prove all the formulae derivable in the ontology as well as their negations. Thus, when trying to diagnose the source of the inconsistency, all the formulae of the ontology and their negations are returned. This is not very informative. As shown above, the use of MUPS and MIS provides many candidate sources of a given contradiction, ultimately requiring a human to decide which repair to effect. As opposed to this, the case of multiple ontologies allows to diagnose the type or shape of the contradiction that arises at the global level once contexts are bridged². Given that the individual ontologies are locally consistent, i.e. reliable, the diagnosis can be focused on *statements that follow from the two ontologies and that directly contradict each other*. This in turn opens up the new possibility of many kinds of syntactical manipulations (e.g., splitting a function, changing its arity, etc.) to re-establish global consistency. Note that a DL-approach like [Ji *et al*, 2009], although focussing on the multiple-ontologies case, does not exploit the potential inherent in it and limits itself to retracting the undesirable mappings. On the other hand, there are earlier non DL-based attempts at automated ontology evolution with multiple-ontologies, such as [McNeill & Bundy, 2007] which investigate how agents with different ontologies interact.

For what concerns the second and third limitation found in the literature, i.e., the inherent lack of generality of DL-based approaches, the GALILEO system bases automated ontology evolution on the use of higher-order logic. Diagnosis and repairs more sophisticated than retracting axioms (e.g., splitting a function into parts, adding new arguments to a function, etc.) require a meta-logic more expressive than DL, or FOL. Modelling ORPs in HOL allows for the existential quantification over ontologies and functions, which is useful in domains containing many ontologies/theories, e.g., natural sciences and general real-world semantics. Moreover, the polymorphism of symbols employed in ORPs permits their high generality. HOL formulae can abstract

²Note that although this multiple-ontologies-approach cannot directly deal with a single inconsistent ontology that is self-contained, it can be integrated by an approach for single ontologies, such as [Haase *et al*, 2005], which exploits the notion of connectedness to isolate the source of the contradiction.

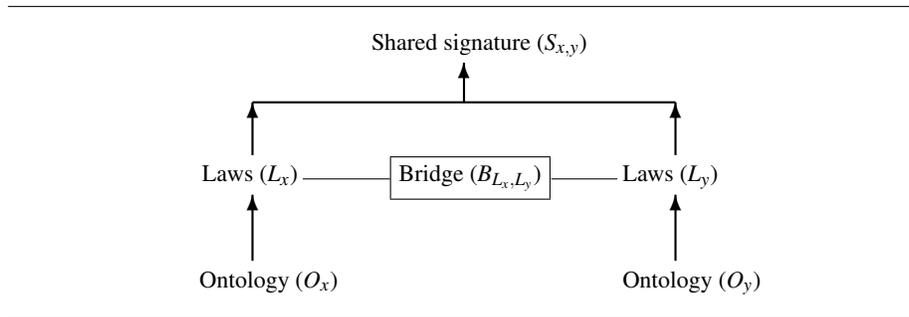


Figure 3: Typical modularization with bridge. O_x , O_y contain data or value assertions about entities in the domain. The arrows pointing upward represent the dependence of O_x and O_y on ontologies L_x resp. L_y , which in turn may share part of their signatures but differ in other parts. B_{L_x, L_y} merges L_x and L_y in one unified context which resolves all the relevant differences by means of so-called bridging axioms.

over types, number of arguments, etc. therefore using HOL as a meta-logic for ontology evolution is a contribution to ontology evolution research. Being able to quantify over the predicates and the ontologies allows for the formulation of logical theories of ontology evolution – as opposed to the algorithms usually found in the DL literature.

Finally, an additional consideration for preferring HOL over DL is its expressivity as an object language, i.e., the possibility of representing directly, rather than as roles in concepts, physics formulae or other inherently higher-order logical parts of mathematics, e.g. calculus.

3 Robustness by ontology repair plans and contexts

At the heart of our approach to ontology evolution lies the notion of ontology repair plan. These are generic combinations of diagnosis and repair operations that guide the evolution of an ontology. The diagnostic component takes as input two or more ontologies that represent physics theories and experimental data, and checks whether specific logical conditions (e.g. a contradiction) hold between the ontologies. In the positive, the repair component is triggered. By grouping these meta-level operations, ORPs tradeoff completeness against reduction in search.

ORPs operate on ontologies that are formalized and implemented as contexts and that are assumed to have modularization patterns akin to the DL distinction between TBox and ABox. As shown in Figure 3, we assume that any two physics ontologies O_x , O_y contain data or value assertions about entities in the domain. The arrows pointing upward represent the fact that O_x and O_y depend on, i.e., inherit all contents of ontologies L_x resp. L_y . These two ontologies contain specifications of the laws of physics, which allow the derivation of new information about the entities from the data in O_x and O_y . L_x and L_y have signatures that are at least in part different, i.e., they use different languages. In order to resolve such differences, a bridging ontology B_{L_x, L_y}

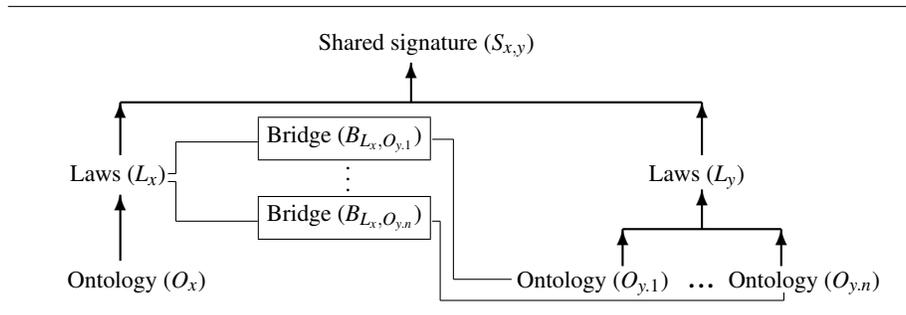


Figure 4: Another typical modularization with bridges. Ontologies $O_{y,1} \dots O_{y,n}$ provide data for L_y and conflicts between such data are detected and resolved by bridging each $O_{y,1} \dots O_{y,n}$ to L_x .

merges in one unified context all symbols of the languages of L_x and L_y and resolves all relevant differences between L_x and L_y by means of so-called bridging axioms. Using information of O_x in B_{L_x, L_y} it is possible to derive theorems in the language of O_y and vice versa. Therefore, if O_x and O_y contain conflicting information about the entities in the domain, the contradiction becomes apparent through the bridging ontology. L_x and L_y do not necessarily share signature elements in an ontology $S_{x,y}$. But, when this is the case, they usually share basic mathematical theories used in physics, e.g., arithmetic or geometry. However, $S_{x,y}$ is not required for ORPs to operate, as long as all representational differences are resolved in the bridging ontology. Figure 4 shows a modularization in which ontologies $O_{y,1} \dots O_{y,n}$ provide data for L_y and conflicts between such data are detected and resolved by bridging each $O_{y,1} \dots O_{y,n}$ to L_x . Earlier presentations of ORPs without bridges and of their applications can be found in [Bundy & Chan, 2008] and [Chan & Bundy, 2008].

So far eight ORPs have been conceptualised and formalised, three of them are being implemented and are being tested.

The ORP called *Where's My Stuff?* (WMS) is triggered in situations such as Figure 3, where a theoretical prediction in O_x conflicts with sensory information derived from experiments in O_y through the bridge B_{L_x, L_y} . WMS deploys an addition-strategy that is quite common in physics. For instance, in order to account for unpredictable yet observed gravitational behaviours in the orbit of a planet or in the stellar orbital velocity in a galaxy, astronomers postulate the presence of an additional unobserved planet or, resp., of dark matter. Accordingly, WMS redefines the contradictory function (in the examples, the functions orbit, resp., orbital velocity) as the sum of a *visible* part (i.e. the amount calculated by the original function) and an *invisible* part (i.e. the amount that can only indirectly be observed).

The ORP called *Inconstancy* is triggered in situations such as Figure 4, where sensory information is derived from experiments run under different circumstances $O_{y,1} \dots O_{y,n}$. If such experiments present variations (e.g. variations in temperature) in the value of a function *stuff* which according to O_x , should be constant (e.g. the volume to pressure ratio of a gas) *Inconstancy* repairs the theoretical ontology by changing the

signature and making the constant quantity dependent on the varying quantity.

The ORP called *Unite* is triggered in situations such as Figure 3, where the observations in O_y for two distinct entities in L_x fully match, and this is a ground for identifying the two functions as one. *Unite* is the converse of WMS: it is not triggered by a contradiction, but by an equality of two *stuff*s relative to a chosen defining property; the repair is to equate in O_x the two *stuff*s. Note that this ORP can be considered as an inconsistency-resolution mechanism to the same extent to which redundancy can be considered as a form of scientific inconsistency, as it violates the general principles of parsimony such as Occam’s razor.

Other ORPs perform other types of repairs, such as merging two theoretical ontologies; changing the type structure of a function to let it fit a given dataset; drawing an analogy between two theoretical ontologies; revising the geometry of a physical object type; skolemizing a theorem containing an existential quantifier and transforming a set of unary predicates into a unary function.

4 The *Where’s My Stuff?* ORP

The *Where’s My Stuff?* ORP, described in Figure 5, is triggered when the predicted value returned by a function conflicts with the observed value of the same function. It assumes three ontologies: an ontology O_x representing the current state of a predictive physics theory, a heterogeneous ontology O_y representing some sensory information arising from an experiment, a bridge B_{L_x,L_y} mapping O_x ’s and O_y ’s heterogeneous signature elements. Suppose the function f measures some property of *stuff*. There are two possible ways to identifying the conflict: with equation (16), where $f(\textit{stuff})$ is equal or less than a particular value in $Ax(B_{L_x,L_y}) \cup Ax(O_y)$ (where $Ax(O)$ returns all axioms of O), and with equation (17), where $f(\textit{stuff})$ is equal or greater than a particular value in $Ax(B_{L_x,L_y}) \cup Ax(O_x)$. The repair is to split *stuff* into three: visible *stuff*, invisible *stuff*, and total *stuff*, defining invisible *stuff* in terms of total and visible *stuff*s in the repaired O_x , $v(O_x)$ (18, 19). The new B_{L_x,L_y} , $v(B_{L_x,L_y})$, is the same as B_{L_x,L_y} except for the renaming of *stuff* to *stuff_{vis}* (20).

4.1 WMS’s application to the discovery of latent heat

Until the second half of the 18th century, the chemical/physical notion of heat was conflated with the notion of temperature and it was seen as a function of a temporal quantity, e.g. flow [Wiser & Carey, 1983]. Flow was the cause of the change of temperature of a physical body put in direct contact with another physical body at a different temperature. This pre-modern view of heat and temperature can be rationally reconstructed in the following equation:

$$\Delta T \equiv Q = m \times \Delta t \quad (21)$$

where Q is the heat absorbed or released by a body and is equivalent to ΔT , the body’s difference in temperature at the start and at the end of the flow, m is the body’s mass, Δt is the flow of heat, measured over time, from the hotter to the cooler body.

In the period 1759-1763 the discoveries of specific and of latent heat by Joseph Black

Trigger: O_x is the predictive ontology and O_y is the observational one. If $f(stuff)$ has two different values in O_x and in $Ax(B_{L_x, L_y}) \cup O_y$ then the following formula will be triggered:

$$\exists L_x, L_y, O_x, O_y, B_{L_x, L_y}: \text{Onto}, \tau, \tau': \text{Types}, f: \tau \mapsto \tau', \text{stuff}: \tau, v: \tau'. \quad (15)$$

$$(L_x \subseteq O_x \wedge L_y \subseteq O_y \rightarrow \quad (15)$$

$$((Ax(O_x) \vdash f(stuff) >_{\tau'} v \wedge Ax(B_{L_x, L_y}) \cup Ax(O_y) \vdash f(stuff) \leq_{\tau'} v) \vee \quad (16)$$

$$(Ax(B_{L_x, L_y}) \cup Ax(O_x) \vdash f(stuff) \geq_{\tau'} v \wedge Ax(O_y) \vdash f(stuff) <_{\tau'} v))) \quad (17)$$

where $L \subseteq O$ means that ontology O depends on ontology L ; B_{L_x, L_y} means that ontology B bridges ontologies L_x and L_y ; $O \vdash \phi$ means that formula ϕ is a theorem of ontology O ; $t: \text{Types}$ means t is a type; $o: \text{Onto}$ means o is an ontology; $>_{\tau}$ is the greater-than operator for τ ; $Ax(O)$ returns all axioms of O . In the case where each comparison operator in (16, 17) is reversed, the roles of O_x and B_{L_x, L_y} below are reversed.

Repair: Two new kinds of stuff and a definition of the *invisible* stuff are introduced.

$$\text{stuff} \sigma_{invis} := \text{stuff} -_{\tau} \text{stuff} \sigma_{vis} \quad (18)$$

When $stuff$ is a constant, the substitution σ_{vis} just replaces it with new constant standing for the visible stuff; when $stuff$ is compound, the replacement is more complex, but still automatable. Similar remarks hold for σ_{invis} .

Let $v(O_x)$ and $v(B_{L_x, L_y})$ be the repaired ontology, resp., bridge. The axioms for the new ontologies are updated in terms of those of the old as follows:

$$Ax(v(O_x)) := \{ \text{stuff} \sigma_{invis} := \text{stuff} -_{\tau} \text{stuff} \sigma_{vis} \} \cup Ax(O_x) \quad (19)$$

$$Ax(v(B_{L_x, L_y})) := \{ \phi \{ \text{stuff} / \text{stuff} \sigma_{vis} \} \mid \phi \in Ax(B_{L_x, L_y}) \} \quad (20)$$

To effect the repair, the axioms of $v(O_x)$ are the same as those of O_x except for the addition of the new definition; the axioms of $v(B_{L_x, L_y})$ are the same as those of B_{L_x, L_y} except for the renaming of the original stuff to the visible stuff.

Figure 5: The *Where's My Stuff?* ontology repair plan with bridge

brought to a revision of the view underlying equation (21) and established the modern view, based on a distinction between the quantity of heat transferred to a physical body and the rise in temperature undergone by the body. On the one hand, the notion of specific heat capacity (1759) accounted for the fact that when equal masses of different materials at equal temperatures absorb the same quantity of heat (i.e. when exposed to equal flows) they undergo different rises in temperature. This is due to a material specific capacity constant c . On the other hand, the case of single-substance bodies led to the formulation of the theory of latent heat (1761). A melting block of, for instance, ice releases heat at constant temperature. Such heat is proportional to a constant L which is specific to the body's material as well as to its phase-transition. This modern view of heat and temperature is represented by the following equation:

$$Q = m \times \Delta T \times c + m \times L \quad (22)$$

$$Ax(S_{1,2}) ::= \{ \dots, Mass::Obj \mapsto Evt \mapsto Real, Melting::Evt, Start::Evt \mapsto Real, \dots \} \quad (23)$$

$$Ax(L_1) ::= \{ \dots, Duration::Evt \mapsto Time, Heat::Obj \mapsto Evt \mapsto Real, \quad (24)$$

$$\forall \alpha:Obj, e:Evt. Heat(\alpha, e) = Mass(\alpha, e) \times Duration(e) \} \quad (25)$$

$$Ax(L_2) ::= \{ \dots, Hght::Obj \mapsto Time \mapsto Real, TempDiff::Obj \mapsto Evt \mapsto Real, \dots \quad (26)$$

$$\forall \alpha:Obj, e:Evt. TempDiff(\alpha, e) = (Hght(MerOf(ThmIn(\alpha, Start(e)))) - Hght(MerOf(ThmIn(\alpha, End(e)))))) \} \quad (27)$$

$$Ax(O_1) ::= \{ Mass(H_2O, Melting) > 0, \quad (28)$$

$$Duration(Melting) > 0 \} \quad (29)$$

$$Ax(O_2) ::= \{ Hght(MerOf(ThmIn(H_2O, Start(Melting)))) = 5, \quad (30)$$

$$Hght(MerOf(ThmIn(H_2O, End(Melting)))) = 5 \} \quad (31)$$

$$Ax(B_{L_1, L_2}) ::= \{ \dots, L_1.Start \rightsquigarrow Start, L_2.Start \rightsquigarrow Start, L_1.Duration \rightsquigarrow Duration, L_2.Hght \rightsquigarrow Hght, \dots \quad (32)$$

$$\forall \alpha:Obj, e:Evt, \forall v:Real. Heat(\alpha, e) = v \iff TempDiff(\alpha, e) = v \} \quad (33)$$

Figure 6: Initial model of the Latent Heat case study. Bridge B_{L_1, L_2} merges (indicated by \rightsquigarrow) L_1 's and L_2 's signatures and maps their concepts, for instance, *Heat*, the amount of heat absorbed or released by a body, is mapped on *TempDiff*, the temperature difference undergone by the body. This allows the contradiction between B_{L_1, L_2} and O_1 about the value of *Heat* to become explicit, just as in the WMS trigger formula, because during a *melting* event there is no change in temperature.

where Q is the heat put into or taken out of the body, m is the mass of the body, c is the specific heat capacity of the body's substance, ΔT is the change in temperature, L is the specific latent heat of the substance during the considered phase-transition.

Note that the evolution from equation (21) to equation (22) marks a shift in the meaning of the variables: Q is not mapped anymore on and measured solely by the temperature difference of the body, ΔT (which now has its own place in the equation). Also, there is no need anymore to make explicit reference to a temporal notion such as flow (Δt). Based on the modularization shown in Figure 3, Figure 6 represents the state of the theory of heat according to equation (21) and to the experimental results that lead to the introduction of latent heat. At the top, $S_{1,2}$ contains the shared signature of two ontologies of physics. For instance, *Melting* is defined as a subtype of event (*Evt*), *Start* as a function from *Evt* to the type of the real numbers *Real*. L_1 introduces new signature elements such as the duration of an event (*Duration*) and represents equation (21) by equation (25). O_1 contains value assertions for *Mass* and *Duration*

that allow this module to predict that the quantity of *Heat* released by H_2O during the *Melting* event is greater than zero. On the other hand, L_2 represents the experimental set-up: it introduces other signature elements for the instruments (such as *ThmIn* for the thermometer used in the measurements) and it defines the difference in temperature of a body (*TempDiff*) as the difference in height (*Hght*) of the mercury column (*MerOf*) of a thermometer stuck in the body throughout an event. O_2 contains value assertions for the relevant observations of the height of the mercury. Bridge B_{L_1,L_2} translates L_1 's and L_2 's signatures. For instance, *Start* in L_1 (indicated by $L_1.Start$) is translated (indicated by the symbol \rightsquigarrow) to *Start*. B_{L_1,L_2} also aligns signature elements, for instance *Heat* and *TempDiff*. This allows the contradiction with O_1 to become explicit just like in the WMS trigger formula (16):

$$Ax(O_1) \vdash Heat(H_2O, Melting) > 0 \quad (34)$$

$$Ax(B_{L_1,L_2}) \cup Ax(O_2) \vdash Heat(H_2O, Melting) = 0 \quad (35)$$

Given the substitution:

$$\{O_1/O_x, O_2/O_y, \lambda x.x(H_2O, Melting)/f, Heat/stuff, 0/v\}$$

WMS repairs the two ontologies as in $v(O_x)$ (i.e. by adding to O_1 the distinction between visible and invisible *Heat*) and as in $v(B_{L_x,L_y})$ (i.e. by renaming in B_{L_1,L_2} all occurrences of heat to *Heat_{vis}*).

$$Ax(v(O_1)) := \{ Heat_{invis} := Heat - Heat_{vis} \} \cup Ax(O_1) \quad (36)$$

$$Ax(v(B_{L_1,L_2})) := \{ \phi \{ Heat/Heat_{vis} \} \mid \phi \in Ax(B_{L_1,L_2}) \} \quad (37)$$

4.2 WMS's application to the postulation of dark matter

The theoretical existence of dark matter is based on various sources of evidence, including the rotational velocities of stars in spiral galaxies, which exceeds the predicted orbital velocities³, as first observed by Rubin [Rubin *et al*, 1980]. Given the observed distribution of mass in these galaxies, Newtonian dynamics predicts that orbital velocities decrease inversely with the square root of the distance from the galactic centre, or the *radius* (equation (41)). However, the observed velocity was almost constant out to large radii. Rubin's conclusion was that some invisible matter exerts a gravitational force on these stars, causing the unexpectedly high orbital velocities. The theoretical notions that play a role in this case are Newton's second law of motion (equation (38)), the law of gravitational attraction in circular orbits (equation (39)) and the law of centripetal acceleration (equation (40)), which combined allow to derive the orbital

³It is assumed that *observed* orbital velocity and rotational velocity for stars are the same. Observations were initially based on gas clouds (HII regions) rather than stars, and only later were the observations based on both gases and stars. We simplify the physics by considering only stars.

velocity of a body (equation (41)) at distance r from the body at the centre of the orbit.

$$F = ma \quad (38)$$

$$F = \frac{GMm}{r^2} \quad (39)$$

$$a = \frac{v_{orb}^2}{r} \quad (40)$$

$$v_{orb} = \sqrt{\frac{G \times M}{r}} \quad (41)$$

where F is the force applied to an orbiting body, m is the mass of the body, a is its acceleration, M is the mass of the body at the centre of the orbit, G is the gravitational constant, r is the radius, or distance between the orbiting body and the body at the centre of the orbit, v_{orb} is the orbital velocity.

On the experimental side, orbital velocity is calculated as rotational velocity, based on spectrographic data and using, among others, equations (42, 43, 44).

$$z = \frac{\lambda - \lambda_0}{\lambda_0} \quad (42)$$

$$v_{rad} = c \times z \quad (43)$$

$$v_{rot} = \frac{v_{rad} - v_{sys}}{\sin(i)} \quad (44)$$

where z is the redshift of a radiation λ with respect to a reference λ_0 , c is the speed of light, v_{rad} is the radial velocity of a body (e.g. a star) along the line of observation, v_{sys} is the velocity of the system (e.g. a galaxy) to which the observed body belongs, v_{rot} is the rotational velocity of the object. Based on the modularization shown in Figure 3, Figure 7 represents the state of the theory of galactic orbital velocity according to equation (41) and to the observations yielded by equations (42) to (44), which led to the postulation of dark matter. At the top, $S_{3,4}$ contains the shared signature where, for instance, the type *Rad* for distances from the galactic centre is defined as a function that maps the product of objects (e.g. stars) and object sets (e.g. galaxies) and an observation event onto the type *Dst* for distances (this in turn maps onto real numbers). Note that $S_{3,4}$ is assumed to inherit from some higher mathematical module knowledge about shapes of curves, e.g., *UpFlat* is the curve that has a positive gradient between zero and some point and a zero gradient thereafter. Such knowledge makes it possible to compare the predictions and the observations. L_3 and O_3 introduce the galaxy rotation curve *GphA*, a Keplerian curve computed according to equation (41) (which models equation (41)) where *Evt*, *Gly*, *Dst*, and *Str* denote types for representing events, galaxies, distances, and stars, resp.; G denotes the universal gravitational constant; *OV*, the orbital velocity; g_r , the set of stars in the galaxy g up to distance r ; *Mass*, the mass of a body. L_4 and O_4 introduce *GphB*, the curve based on observations, computed according to equations (50) to (52) (which model equations (42) to (44)) where *RtV* denotes the rotational velocity; *RdV*, the radial velocity; *SyV*, the velocity of the galactic system relative to the observer; *Inc*, the inclination of the galaxy; *λShift*,

$$Ax(S_{3,4}) ::= \{ \dots, UpFlat :: Real \mapsto Real, Rad :: (Obj * Obj Set) \mapsto Evt \mapsto Dst \dots \} \quad (45)$$

$$Ax(L_3) ::= \{ \dots, GphA :: Gly \mapsto Evt \mapsto Real, \dots \} \quad (46)$$

$$\forall eEvt, gGly, rDst. OV(g, e) = \sqrt{\frac{G \times \sum s \in g_r.Mass(s, e)}{r}}, \quad (47)$$

$$\forall eEvt. GphA(Glxy71, e) = OV(Glxy71, e) \quad (48)$$

$$Ax(L_4) ::= \{ \dots, GphB :: Gly \mapsto Evt \mapsto Real, \dots \} \quad (49)$$

$$\forall eEvt, sStr, gGly. RdV(s, g, e) = c \times \lambda Shift(s, g, e), \quad (50)$$

$$\forall eEvt, sStr, gGly. RtV(s, g, e) = \frac{RdV(s, g, e) - SyV(g, e)}{\sin(Inc(g, e))}, \quad (51)$$

$$\forall sStr. GphB(Glxy71, Rad(s, Glxy71)) = RtV(s, Glxy71, Obs6) \quad (52)$$

$$Ax(O_3) ::= \{ G = 6.673 \times 10^{-11}, \quad (53)$$

$$\sum s \in Glxy71_{Rad(Star1, Glxy71)}.Mass(s, Obs6) = 100, \quad (54)$$

$$\sum s \in Glxy71_{Rad(Star9, Glxy71)}.Mass(s, Obs6) = 110 \} \quad (55)$$

$$Ax(O_4) ::= \{ c = 299792458, \quad (56)$$

$$\lambda Shift(Star1, Glxy71, Obs6) = 300, \quad (57)$$

$$\lambda Shift(Star9, Glxy71, Obs6) = 300 \} \quad (58)$$

$$Ax(B_{L_3, L_4}) ::= \{ \dots, L_3.UpFlat \rightsquigarrow UpFlat, L_4.UpFlat \rightsquigarrow UpFlat, \dots \} \quad (59)$$

$$\forall eEvt, sStr, gGly, v, rReal. RtV(s, g, e) = v \wedge Rad(s, g) = r \iff OV(g, r, e) = v \} \quad (60)$$

Figure 7: Initial model of the Dark Matter case study. B_{L_3, L_4} merges (indicated by \rightsquigarrow) the signatures of L_3 and L_4 and aligns RtV with OV . This allows the contradiction with O_4 to become explicit just as in the reversed WMS trigger formula (17).

the shift in wave length; $Glxy71$, the galaxy being observed; $Star1$ and $Star9$, stars in the observed galaxy; and, $Obs6$, the observation event.

Some symbols in O_4 , such as RtV and RdV , are not in the language of O_3 . B_{L_3, L_4} links together the seemingly disparate terms by relating RtV in L_4 to the OV in L_3 . This allows the contradiction with O_4 to become explicit just like in the reversed WMS trigger formula (17):

$$Ax(O_3) \vdash GphA < UpFlat \quad (61)$$

$$Ax(B_{L_3, L_4}) \cup Ax(O_4) \vdash GphA = UpFlat \quad (62)$$

Given the substitution:

$$\{O_3/O_x, O_4/O_y, GphA/f, Glxy71/stuff, UpFlat/v\}$$

WMS repairs the two ontologies as in $v(O_x)$ (i.e. by adding to O_3 the distinction between visible and invisible $Glxy71$) and as in $v(B_{L_x,L_y})$ (i.e. by renaming in B_{L_3,L_4} all occurrences of $Glxy71$ to $Glxy71_{vis}$).

$$Ax(v(O_3)) := \{ \phi \{ Glxy71 / Glxy71_{vis} \} \mid \phi \in Ax(O_3) \} \quad (63)$$

$$Ax(v(B_{L_3,L_4})) := \{ Glxy71_{invis} := Glxy71 - Glxy71_{vis} \} \cup Ax(B_{L_3,L_4}) \quad (64)$$

5 The *Inconstancy* ORP

The *Inconstancy* ORP, described in Figure 8, is triggered when, given an ontology O_x representing the current state of a physical theory and some ontologies $O_{y,i}$ representing sensory information arising from experiments, the sensory ontologies give distinct values for function $stuff(\vec{s}_i)$ in different circumstances. Suppose function $V(\vec{s}_i, \vec{b}_i)$ of the i^{th} sensory ontology, where \vec{b}_i contains variables distinguishing among these circumstances, returns distinct values in each of these circumstances, but is *not* one of the parameters in \vec{s}_i , i.e., $stuff(\vec{s}_i)$ does not depend on $V(\vec{s}_i, \vec{b}_i)$. We call $stuff(\vec{s}_i)$ the *inconstancy* and $V(\vec{s}_i, \vec{b}_i)$ the *variad*. The *Inconstancy* repair plan establishes a relationship between the variad $V(\vec{s}_i, \vec{b}_i)$ and the inconstancy $stuff(\vec{s}_i)$. The inconstancy might, for instance, be the gravitational constant G and the variad might be the acceleration of an orbiting star due to gravity, as suggested by the MODified Newtonian Dynamics approach to the dark matter case study.

To discover the meaning of the function F , *Inconstancy* follows the tradition of Langley's BACON program [Langley *et al.*, 1983] by using curve fitting. The ontologies $O_s(V(\vec{s}_i, \vec{b}_i) = v_i \dots)$ provide a useful collection of equations: $F(c(\vec{s}_i), V(\vec{s}_i, \vec{b}_i)) = c_i$ for $i = 1, \dots, n$. Curve fitting techniques, e.g., regression analysis, can be applied to these equations to approximate a definition of F . This hypothesis can then be tested by creating additional observations $O_s(V(\vec{s}_j, \vec{b}_j) = v_j \dots)$, for new values of $V(\vec{s}_j, \vec{b}_j)$, and confirming or refuting the hypothesis.

5.1 *Inconstancy* application to the MOND approach

An alternative theory for the anomaly in orbital velocities of stars in galaxies is provided by MODified Newtonian Dynamics MOND, proposed by Moti Milgrom in 1981 as an alternative to the dark matter explanation. MOND is an example of the *Inconstancy* plan. This is a good example of how the same observational discrepancies can trigger different repair plans. MOND suggests that the gravitational constant is not a constant, but depends on the acceleration between the objects on which it is acting. It is constant until the acceleration becomes very small and then it depends on this acceleration, which is the case for stars in spiral galaxies. So, the gravitational constant G in equation (41) can be repaired by giving it an additional argument to become $G(a)$, where a is the centripetal acceleration of a star s (equation (40)) due to the gravitational attraction between the star and the galaxy in which it belongs. a is the variad and G is the inconstancy.

Based on the modularization shown in Figure 4, Figure 9 extends and modifies the axiomatization presented in Figure 7. $S_{5,6}$ extends $S_{3,4}$ to include the symbol for cen-

Trigger: O_x is the predictive ontology, $O_{y,i}(v(\vec{b}_i) =_{\tau'} v_i) \forall i \in [1, n]$ are the observational ones, made under the conditions $v(\vec{b}_i) =_{\tau'} v_i$. $O_{y,i}$ are bridged to O_x by $B_{L_x, O_{y,i}}$. If *stuff* is measured to take different values in at least two of these bridges, the following formulae will be triggered:

$$\begin{aligned} \exists L_x, L_y, O_x, O_{y,1}, O_{y,n}, B_{L_x, O_{y,1}}, B_{L_x, O_{y,n}} : \text{Onto}, \tau, \tau', \vec{\tau}'' : \text{Types}, \text{stuff}, c, c_1, \dots, c_n : \tau, \\ v_1, \dots, v_n : \tau', \vec{b}_1, \dots, \vec{b}_n : \vec{\tau}''; v : \vec{\tau}'' \mapsto \tau' \\ L_x \subseteq O_x \wedge L_y \subseteq O_y \rightarrow \end{aligned} \quad (65)$$

$$(Ax(O_x) \vdash f(\text{stuff})) ::= c \wedge \quad (66)$$

$$\begin{aligned} Ax(B_{L_x, O_{y,1}}(v(\vec{b}_1) =_{\tau'} v_1)) \vdash f(\text{stuff}) =_{\tau} c_1 \wedge \\ \vdots \\ Ax(B_{L_x, O_{y,n}}(v(\vec{b}_n) =_{\tau'} v_n)) \vdash f(\text{stuff}) =_{\tau} c_n \wedge \\ Ax(O_x) \vdash c_1 \neq c_n \end{aligned} \quad (67)$$

$$Ax(O_x) \vdash c_1 \neq c_n \quad (68)$$

where $L \subseteq O$ means that ontology O depends on ontology L ; B_{L_x, L_y} means that ontology B bridges ontologies L_x and L_y ; $O \vdash \phi$ means that formula ϕ is a theorem of ontology O ; $t : \text{Types}$ means t is a type; $o : \text{Onto}$ means o is an ontology; $>_{\tau}$ is the greater-than operator for τ ; $Ax(O)$ returns all axioms of O .

Repair: The repair is to change the signature of all the ontologies to relate the inconstancy, *stuff*, to the variad, $v(\vec{y})$ via a new function F :

$$v(\text{stuff}) ::= \lambda \vec{y}. F(c, v(\vec{y})) \quad (69)$$

The axioms of the new ontologies are calculated in terms of those of the old, as follows:

$$\begin{aligned} Ax(v(O_x)) ::= \{ \phi \{ \text{stuff} / v(\text{stuff})(\vec{y}) \} \mid \phi \in Ax(O_x) \} \\ \setminus \{ \text{stuff} ::= c \} \cup \{ v(\text{stuff}) ::= \lambda \vec{y}. F(c, v(\vec{y})) \} \end{aligned} \quad (70)$$

$$\begin{aligned} Ax(v(B_{L_x, O_{y,i}}(v(\vec{b}_i) =_{\tau'} v_i))) ::= \\ \{ \phi \{ \text{stuff} / v(\text{stuff})(\vec{b}_i) \} \mid \phi \in Ax(B_{L_x, O_{y,i}}(v(\vec{b}_i) =_{\tau'} v_i)) \} \end{aligned} \quad (71)$$

Figure 8: The *Inconstancy* ontology repair plan with bridges

tripetal acceleration $Accel(s, e)$. This is defined in L_3 's extension L_5 as a function of orbital velocity and in L_4 's extension L_6 as a function of rotational velocity. $O_{6,1}$ and $O_{6,2}$ contain the data that allow to calculate according to equation (51) the rotational velocity for *Star1* and *Star9*, which has the same value. Therefore, according to equation (74) their acceleration varies with their radius. On the other had, bridges $B_{L_5, O_{6,1}}$ and $B_{L_5, O_{6,2}}$, which just as B_{L_3, L_4} , equate orbital and rotational velocity, allow to calculate the local value for G , which varies as the variad. This allows the contradiction with O_5 to become explicit, like in *Inconstancy*'s trigger formulae (66) through (68), where G is calculate in each $B_{L_5, O_{6,i}}$ by the inverse of equation (47) (the value of which

$$Ax(S_{5,6}) ::= Ax(S_{3,4}) \cup \{ \dots, Accel :: Obj \mapsto Evt \mapsto Real \dots \} \quad (72)$$

$$Ax(L_5) ::= Ax(L_3) \cup \{ \forall e Evt, s Str, g Gly. Accel(s, e) = \frac{(OV(s, g, e))^2}{Rad(s, g)} \} \quad (73)$$

$$Ax(L_6) ::= Ax(L_4) \cup \{ \forall e Evt, s Str, g Gly. Accel(s, e) = \frac{(RtV(s, g, e))^2}{Rad(s, g)} \} \quad (74)$$

$$Ax(O_5) ::= Ax(O_3) \quad (75)$$

$$Ax(O_{6,1}) ::= \{ c = 299792458, \quad (76)$$

$$\lambda Shift(Star1, Glxy71, Obs6) = 300 \} \quad (77)$$

$$Ax(O_{6,9}) ::= \{ c = 299792458, \quad (78)$$

$$\lambda Shift(Star9, Glxy71, Obs6) = 300 \} \quad (79)$$

$$Ax(B_{L_5, O_{6,1}}) ::= Ax(B_{L_3, L_4}) \quad (80)$$

$$Ax(B_{L_5, O_{6,9}}) ::= Ax(B_{L_3, L_4}) \quad (81)$$

Figure 9: Initial model of the MOND case study. Bridges $B_{L_5, O_{6,1}}$ and $B_{L_5, O_{6,2}}$, which just as B_{L_3, L_4} in Figure 7, equate orbital and rotational velocity, allow to calculate the local value for G , which varies as the variad. This allows the contradiction with O_5 to become explicit.

is equated here to G_i for the sake brevity in the substitutions that follow):

$$Ax(O_5) \vdash G ::= 6.673 \times 10^{-11} \quad (82)$$

$$Ax(B_{L_5, O_{6,1}}(Accel(Star1) = A_1)) \vdash$$

$$G = \frac{(OV(Star1, Glxy71, Obs6))^2 \times Rad(Star1, Glxy71)}{\sum s \in Glxy71 / Star1.Mass(s, Obs6)} = G_1 \quad (83)$$

$$\vdots \quad \vdots$$

$$Ax(B_{L_5, O_{6,9}}(Accel(Star9) = A_2)) \vdash$$

$$G = \frac{(OV(Star9, Glxy71, Obs6))^2 \times Rad(Star9, Glxy71)}{\sum s \in Glxy71 / Star9.Mass(s, Obs6)} = G_9 \quad (84)$$

$$Ax(O_5) \vdash G_1 \neq G_9 \quad (85)$$

Given the substitutions:

$$\forall i \in [1, n]$$

$$\{O_5/O_x, O_{6_i}/O_{y_i}, stuff/G, =/f, c/6.673 \times 10^{-11}, c_i/G_i, v/Acc, \vec{b}_i/\langle Star_i \rangle, v_i/A_i\}$$

Inconstancy repairs the two ontologies by redefining the inconstancy G as a function of G and of the variad *Accel* and by replacing the old definition with the new one:

$$\begin{aligned} Ax(v(O_5)) &::= \{\phi\{G/v(G)(\langle s \rangle)\} \mid \phi \in Ax(O_5)\} \\ \setminus \{G &::= 6.673 \times 10^{-11}\} \cup \{v(G) &::= \lambda\langle s \rangle. F(G, Accel(\langle s \rangle))\} \end{aligned} \quad (86)$$

$$\begin{aligned} Ax(v(B_{L_5, O_{6_i}}(Accel(\langle Star_i \rangle) =_{\tau} A_i))) &::= \\ \{\phi\{G/v(G)(\langle Star_i \rangle)\} \mid \phi \in Ax(B_{L_x, O_{y_i}}(Accel(\langle Star_i \rangle) =_{\tau} A_i))\} \end{aligned} \quad (87)$$

6 The *Unite* ORP

The *Unite* ORP, described in Figure 10, is triggered when the observed values of the defining properties of two distinct functions are the same. For instance, the orbit of (heavenly) bodies is one of their defining properties, i.e., equating two orbits is equivalent to identifying the corresponding bodies as the same one, according to principle that two objects cannot be at the same place at the same time. Just like WMS, *Unite* assumes three ontologies: an ontology O_x representing a predictive physics theory and labelling a given property as defining, a heterogeneous ontology O_y representing some sensory information arising from observations, a bridge B_{L_x, L_y} . Contrary to WMS, *Unite* is not triggered by a contradiction, but by an equality of two *stuff*'s relative to the chosen defining quality. The repair is to equate in O_x *stuff*₁ and *stuff*₂.

6.1 *Unite* application to the Morning and Evening Star case study

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. 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).

One way of comparing the two orbits is to calculate their *mean anomaly* M , i.e., the parameter relating position and time for a body moving in a Kepler orbit:

$$M = \sqrt{\frac{G(M+m)}{a^3}} \times t \quad (94)$$

where a is the length of the orbit's semi-major axis, M and m are the orbiting masses, and G is the gravitational constant.

For instance the mean anomaly of specific bodies orbiting the Sun can be understood as the time since the last point of closest approach to the Sun (periapsis) multiplied by their mean motion.

Trigger: O_x is the predictive ontology and O_y is the observational one. If $f(stuff_1)$ and $f(stuff_2)$ has no values in O_x , the same values in O_y , and f is labelled as defining in O_x , then the following formula will be triggered:

$$\exists L_x, L_y, O_x, O_y, B_{L_x, L_y}:Onto, \tau, \tau':Types, f:\tau \mapsto \tau', stuff_1, stuff_2:\tau. \quad (88)$$

$$L_x \subseteq O_x \wedge L_y \subseteq O_y \rightarrow \quad (88)$$

$$(Ax(O_x) \not\vdash stuff_1 =_\tau stuff_2) \quad (89)$$

$$Ax(O_x) \vdash DefProp(f, \tau) \quad (90)$$

$$Ax(B_{L_x, L_y}) \cup Ax(O_y) \vdash f(stuff_1) =_{\tau'} f(stuff_2) \quad (91)$$

where $L \subseteq O$ means that ontology O depends on ontology L ; B_{L_x, L_y} means that ontology B bridges ontologies L_x and L_y ; $O \vdash \phi$ means that formula ϕ is a theorem of ontology O ; $t:Types$ means t is a type; $o:Onto$ means o is an ontology; $=_\tau$ and $=_{\tau'}$ are the equality operators for τ and τ' ; $Ax(O)$ returns all axioms of O .

Repair: The two *stuff*'s are identified as one.

$$stuff \sigma_{invis} := stuff_1 =_\tau stuff_2 \quad (92)$$

Let $v(O_x)$ be the repaired predictive ontology. The axioms for the new ontologies are updated in terms of those of the old as follows:

$$Ax(v(O_x)) := \{ stuff_1 =_\tau stuff_2 \} \cup Ax(O_x) \quad (93)$$

To effect the repair, the axioms of $v(O_x)$ are the same as those of O_x except for the addition of the new definition.

Figure 10: The *Unite* ontology repair plan with bridge

Based on the modularization shown in Figure 3, Figure 11 axiomatizes the case study at hand. $S_{9,10}$ provides type declarations for *MorningStar*, *EveningStar*, *Sun*. L_9 provides theoretical knowledge for calculating the mean anomaly of objects orbiting the Sun and declares the defining property; O_9 is empty, no prediction is made using the new heliocentric theory. On the other hand, L_{10} is empty, as the old geocentric theory is not relevant, while the observations of *MorningStar* and *EveningStar* in O_{10} are. $B_{L_9, L_{10}}$ allows to apply the new theory to the available data as follows:

$$Ax(O_9) \not\vdash MorningStar =_\tau EveningStar \quad (107)$$

$$Ax(O_9) \vdash DefProp(MeanAnom(o, e), o) \quad (108)$$

$$Ax(B_{L_9, L_{10}}) \cup Ax(O_{10}) \vdash \forall e. MeanAnom(MorningStar, e) =_{\tau'} MeanAnom(MorningStar, e) \quad (109)$$

Given the substitutions:

$$\{O_9/O_x, O_{10}/O_y, stuff_1/MorningStar, stuff_2/EveningStar, f/\lambda e.(MeanAnom(o, Sun, Moon, e))\}$$

$$Ax(S_{9,10}) ::= \{ \dots, MorningStar::Obj, EveningStar::Obj, Sun::Obj \dots \} \quad (95)$$

$$Ax(L_9) ::= \{ \dots, MeanAnom::Obj \mapsto Evt \mapsto RealSet, \dots \} \quad (96)$$

$$\forall \alpha Obj, e Evt. MeanAnom(o, e) = \sqrt{\frac{G \times (Mass(Sun, e) + Mass(o, e))}{Rad_a^2(o, Sun, e)}} \times Duration(e), \quad (97)$$

$$DefProp(MeanAnom(o, e), o) \quad (98)$$

$$Ax(L_{10}) ::= \{ \} \quad (99)$$

$$Ax(O_9) ::= \{ \} \quad (100)$$

$$Ax(O_{10}) ::= \{ \forall e. Mass(Sun, e) = 1.9891 \times 10^{30}, \quad (101)$$

$$\forall e. Mass(MorningStar, e) = 4.868 \times 5 \times 10^{24}, \quad (102)$$

$$\forall e. Mass(EveningStar, e) = 4.868 \times 5 \times 10^{24}, \quad (103)$$

$$Rad(MorningStar, Obs1) = 107940156, \quad (104)$$

$$Rad(EveningStar, Obs1) = 107940156, \quad (105)$$

⋮

$$Rad(MorningStar, Obs5) = 108942109, \quad (106)$$

$$Rad(EveningStar, Obs5) = 108942109 \quad (107)$$

$$Ax(B_{L_9, L_{10}}) ::= \{ \dots, L_9. MeanAnom \rightsquigarrow MeanAnom, \dots \} \quad (108)$$

Figure 11: Initial model of the Morning and Evening Star case study: $B_{L_9, L_{10}}$ allows to calculate the orbit of *MorningStar* and *EveningStar* from the data available in O_{10} and thus identify the two stars as one heavenly body.

Unite repairs O_9 as follows:

$$Ax(v(O_9)) ::= Ax(O_9) \cup \{ MorningStar =_{\tau} EveningStar \}$$

6.2 *Unite* application to the shape of the Earth case study

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, among which were 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.

The defining property is $Project(v, Sun, Moon, e)$: the orthographic projection of a volume v from the Sun onto the Moon, which for a couple of coordinates x, y is

$$Ax(S_{11,12}) ::= \{ \dots, Shape::Obj \mapsto shape, Earth::Obj, Sun::Obj \dots \} \quad (110)$$

$$Ax(L_{11}) ::= \{ \dots, Project::Obj \mapsto Evt \mapsto RealSet, \dots \} \quad (111)$$

$$\forall o_1, o_2, o_3::Obj, e::Evt. Project(Shape(o_1), o_2, o_3, e) = \quad (112)$$

$$\cos Lat(Shape(o_1), e) \times \sin(Lon(Shape(o_1)) - Lon_0, e) \dots, \\ DefProp(Project(p, o_2, o_3, e), p) \} \quad (113)$$

$$Ax(L_{12}) ::= \{ \}$$

$$Ax(O_{11}) ::= \{ \}$$

$$Ax(O_{12}) ::= \{ \forall e. Shape(Earth, e) = Sphere, \quad (114)$$

$$\forall e. Shape(Ball, e) = Sphere, \quad (115)$$

$$\cos Lat(Shape(Earth), Obs3) \times \sin(Lon(Shape(o_1), Obs3) - Lon_0) = 3, \quad (116)$$

$$\cos Lat(Shape(Ball), Obs3) \times \sin(Lon(Shape(o_1)) - Lon_0, Obs3) = 3, \quad (117)$$

$$\vdots$$

$$\cos Lat(Shape(Earth), Obs4) \times \sin(Lon(Shape(o_1), Obs4) - Lon_0) = 5, \quad (118)$$

$$\cos Lat(Shape(Ball), Obs4) \times \sin(Lon(Shape(o_1)) - Lon_0, Obs4) = 5 \} \quad (119)$$

$$Ax(B_{L_{11}, L_{12}}) ::= \{ \dots, L_{11}.Project \rightsquigarrow Project, \dots \} \quad (120)$$

Figure 12: Initial model of the shape of the Earth case study: $B_{L_{11}, L_{12}}$ allows to calculate the shape of *Earth* and *Ball* from the data available in O_{12} and thus identify the two shapes as the same one.

calculated as follows:

$$x = \cos Lat \times \sin(Lon - Lon_0) \quad (121)$$

$$y = \cos Lat_0 \times \sin Lat - \sin Lat_1 \times \cos Lat \cos(Lon - Lon_0) \quad (122)$$

where *Lat* is the latitude, *Lon* is the longitude, and Lon_0 and Lat_0 are reference longitudes and latitudes, respectively.

The idea is that if two 3D objects always have the same 2D projections then they have the same shape. 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. Note also that the observed projections of *Ball* are a thought experiment.

Based on the modularization shown in Figure 3, Figure 12 axiomatizes the case study at hand. $S_{11,12}$ provides type declarations for, among others, the symbols *Shape*,

Earth, Sun. L_{11} provides part of theoretical knowledge for calculating orthographic projections and declares the defining property for such knowledge, O_{11} is empty, no prediction is made using the sphere theory. On the other hand, L_{12} is empty, as the old (e.g. flat Earth) theory is not relevant while the observations of lunar eclipses in O_{12} are. $B_{L_{11},L_{12}}$ allows to apply the theory to the available data as follows:

$$Ax(O_{11}) \not\vdash Shape(Earth) =_{\tau} Shape(Ball) \quad (123)$$

$$Ax(O_{11}) \vdash DefProp(Project(p, o_2, o_3, e), p) \quad (124)$$

$$Ax(B_{L_{11},L_{12}}) \cup Ax(O_{12}) \vdash \forall e. Project(Shape(Earth), Sun, Moon, e) =_{\tau'} Project(Shape(Ball), Sun, Moon, e) \quad (125)$$

Given the substitutions:

$$\{O_{11}/O_x, O_{12}/O_y, stuff_1/Earth, stuff_2/Ball, f/\lambda e.(Project(Shape(Earth), Sun, Moon, e))\}$$

Unite repairs O_{11} as follows:

$$Ax(\nu(O_{11})) ::= Ax(O_{11}) \cup \{Shape(Earth) =_{\tau} Shape(Ball)\}$$

7 Initial implementation

GALILEO is prototyped in the higher-order theorem prover Isabelle [Paulson, 1994]. On the one hand the ontologies are represented as *Locales* [Ballarin, 2004], i.e., as independent proof contexts that can be extended and merged. Extensions allow to express dependencies between the ontologies (e.g., the arrows in Figures 3 and 4), whereas merges (e.g. the boxes in Figures 3 and 4) allow to align heterogeneous ontologies. On the other hand, ORPs are being implemented as extensions of Isabelle’s higher-order matching (HOM) mechanism, a special case of higher-order unification (HOU). So far we have mainly worked on the implementation of ORPs that diagnose and are triggered by conflicts between ontologies, e.g., WMS and *Inconstancy*.

The goal of conflict diagnosis is the correct instantiation of the term *stuff* in the trigger, and this is essentially a reasoning task. Conflict diagnosis consists of two main phases: preparation, term discovery, heuristics-based filtering.

Preparation requires user interaction. In order to reduce the inferential search scope, the user provides the conflicting ontologies and the bridges (e.g., $O_1, O_2, B_{L_1,B_{L_2}}$) as well as the *shape* of the contradiction, i.e. the derivable sentence causing the contradiction (e.g., $Heat(H_2O, Melting) = 0$).

Term discovery is completely automatic and based on HOM. Due to the high generality of the trigger formulae, where f and *stuff* are polymorphic, the space of possible instantiations typically contains a large number of hits. These, though, often contain arbitrary λ -expressions that have no meaningful semantic relevance.

Heuristics-based filtering is used to prune the solution space, e.g., disregarding instantiations that contain the identity function, or that contain no element of the signature, or that contain the same functional symbol applied to permuted arguments.

Once a match is selected, the repair then takes place. Occurrences of the instantiation of *stuff* in an axiom are checked by HOM as well. Even though the transformation rules of repair require only syntactic analysis of the axioms, formalising repair as a HOM problem provides more robustness to the whole procedure. Thus, the look-up of occurrences of *stuff* in a formula involves searching for a match.

8 Conclusion

Based on an initial discussion of the relevant literature, we have presented a higher-order logical approach to the problem of ontology evolution from the perspective of automating the mechanisms to repair locally consistent but globally inconsistent ontologies. Three ontology repair plans and applications for each them were discussed. The ORP *Where's My Stuff?* was applied to the discovery of latent heat and to the postulation of dark matter; *Inconstancy* was applied to the Modified Newtonian Dynamics approach to the study galaxies; *Unite* was applied to the identification of the Morning and Evening Stars and to the assessment of the shape of the Earth.

Both the HOL formalization and GALILEO's implementation in a higher-order theorem prover are providing us with important elements of evaluation. At this stage we are mainly evaluating the generality of the diagnosis and the meaningfulness of the discovered instantiations and getting the following:

1. The expressivity of higher-order logic is valuable in achieving generality. The instantiations of *stuff* are diverse in types and arities, the described mechanism for diagnosis indeed has the capacity to identify conflicts on a general level – not just within an example, but across disparate case studies.
2. Even at a lower unification bound, the raw search spaces for all case studies are too vast. Therefore the use of heuristics prunes away from the solution scope meaningless solution in the order of the tens of thousands, thus making the search space manageable.
3. Even under ideal circumstances, in which the heuristics prune away all, and only, semantically meaningless matches, the pruned solution space generated typically still contains multiple matches, each corresponding to a unique set of instantiations. The instantiation of *stuff* in each of the matches is logically valid and, hopefully, meaningful. For instance, in the case study of the postulation of dark matter, instantiating *stuff* to *Glxy71* indicates that the galaxy should have an extra component; to *GphA* leads to a redefinition of all predicted curves; and, to *GphA(Glxy71)* limits the modification to only the definition of the predicted curve for *Glxy71*. Similar considerations hold for all other case studies. If the preference for particular matches is indicated, then a manual decision for the preferred matches can be avoided.

4. Compared to the approaches found in the literature and discussed in Section 2, our approach expands the space of solutions. The application of an ontology repair plan often combines the retraction of axioms, the change of beliefs as well the deeper modification of the language in which the ontology is represented. For instance, if one were to apply WMS to the incoherent and inconsistent ontology network shown in Figure 2, WMS would propose an additional solution with respect to DL approaches, that would look as an abstract version of the following:
 - a. add *ToBeTrainedStaff* and *NotToBeTrainedStaff* to signature $sig(\mathcal{T}_{com})$;
 - b. rename *Staff* to *NotToBeTrainedStaff* in \mathcal{T}_{com} ;
 - c. add axiom $Staff \equiv ToBeTrainedStaff \sqcup NotToBeTrainedStaff$ to $sig\mathcal{T}_{com}$;
 - d. rename *Staff* to *ToBeTrainedStaff* in $Map(\mathcal{T}_{uni}^{**}, \mathcal{T}_{com})$.

Note that the abstract solution proposed by WMS would add to $sig(\mathcal{T}_{com})$, for instance, $Staff_{vis}$ rather than *ToBeTrainedStaff* and $Staff_{invis}$ for *NotToBeTrainedStaff*. Yet WMS would automatically change the language of the ontology thus give a very circumscribed hint on where and how the language of the ontology needs to be further refined.

Future work will concentrate on furthering our understanding of the relationships between DL and HOL approaches to ontology evolution; on the logical status of bridges with respect to their implementation as Isabelle locales; on the implementation of other ORPs; on more complex testing and evaluation.

References

- | | |
|-------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| [Ballarin, 2004] | Ballarin, C. (2004). Locales and locale expressions in Isabelle/Isar. <i>Lecture Notes in Computer Science</i> , 3085:34–50. |
| [Bouquet <i>et al</i> , 2004] | Bouquet, P., Giunchiglia, F., Van Harmelen, F., Serafini, L. and Stuckenschmidt, H. (2004). Contextualizing ontologies. <i>Web Semantics: Science, Services and Agents on the World Wide Web</i> , 1(4):325–343. |
| [Bundy & Chan, 2008] | Bundy, Alan and Chan, Michael. (2008). Towards ontology evolution in physics. In Hodges, Wilfrid and de Queiroz, Ruy, (eds.), <i>Logic, Language, Information and Computation</i> , volume 5110 of <i>Lecture Notes in Computer Science</i> , pages 98–110. Springer Berlin / Heidelberg. |
| [Chan & Bundy, 2008] | Chan, M. and Bundy, A. (November 2008). Inconstancy: An ontology repair plan for adding hidden variables. In Bringsjord, S. and Shilliday, A., (eds.), <i>Symposium on Automated Scientific Discovery</i> , number FS-08-03 in Technical Report, pages 10–17. AAAI Press. ISBN 978-1-57735-395-9. |

- [Doan *et al*, 2004] Doan, A., Madhavan, J., Domingos, P. and Halevy, A. (2004). Ontology matching: A machine learning approach. *Handbook on Ontologies in Information Systems*, pages 397–416.
- [Flouris, 2006] Flouris, G. (2006). On Belief Change and Ontology Evolution.
- [Flouris *et al*, 2006] Flouris, G., Z., Huang, Pan, J., Plexousakis, D. and H., Wache. (2006). Inconsistencies, Negations and Changes in Ontologies. pages 1295–1300.
- [Ghilardi *et al*, 2006] Ghilardi, S., Lutz, C. and Wolter, F. (2006). Did I damage my ontology? A case for conservative extensions in description logic. *Proceedings of the Tenth International Conference on Principles of Knowledge Representation and Reasoning*, pages 187–197.
- [Giunchiglia & Shvaiko, 2004] Giunchiglia, F. and Shvaiko, P. (2004). Semantic matching. *The Knowledge Engineering Review*, 18(03):265–280.
- [Haase *et al*, 2005] Haase, P., van Harmelen, F., Huang, Z., Stuckenschmidt, H. and Sure, Y. (2005). A Framework for Handling Inconsistency in Changing Ontologies. In *The Semantic Web-ISWC 2005: 4th International Semantic Web Conference, ISWC 2005, Galway, Ireland, November 6-10, 2005: Proceedings*. Springer Verlag.
- [Ji *et al*, 2009] Ji, Q., Haase, P., Qi, G., Hitzler, P. and Stadtmüller, S. (2009). RaDON Repair and Diagnosis in Ontology Networks. *The Semantic Web: Research and Applications*, pages 863–867.
- [Kalyanpur *et al*, 2006a] Kalyanpur, A., Parsia, B., Sirin, E. and Grau, B.C. (2006a). Repairing unsatisfiable concepts in OWL ontologies. In *ESWC*, pages 170–184. Springer.
- [Kalyanpur *et al*, 2006b] Kalyanpur, A., Parsia, B., Sirin, E., Grau, B.C. and Hendler, J. (2006b). Swoop: A web ontology editing browser. *Web Semantics: Science, Services and Agents on the World Wide Web*, 4(2):144–153.
- [Lam *et al*, 2008] Lam, J.S.C., Sleeman, D., Pan, J.Z. and Vasconcelos, W. (2008). A fine-grained approach to resolving unsatisfiable ontologies. *Lecture Notes in Computer Science*, 4900:62.

- [Langley *et al*, 1983] Langley, P., Zytkow, J., Simon, H. and Bradshaw, G. (June 1983). Mechanisms for qualitative and quantitative discovery. In Michalski, R. S, (ed.), *Proceedings of the International Machine Learning Workshop*, pages 121–132. University of Illinois.
- [McNeill & Bundy, 2007] McNeill, F. and Bundy, A. (2007). Dynamic, automatic, first-order ontology repair by diagnosis of failed plan execution. *International Journal On Semantic Web and Information Systems*, 3(3):1–35. Special issue on ontology matching.
- [Paulson, 1994] Paulson, L. C. (1994). *Isabelle: A generic theorem prover*. Springer-Verlag.
- [Rubin *et al*, 1980] Rubin, V. C., Thonnard, N. and Ford, Jr. W. K. (1980). Rotational properties of 21 SC galaxies with a large range of luminosities and radii, from NGC 4605 (R = 4kpc) to UGC 2885 (R = 122 kpc). *Astrophysical Journal*, 238:471.
- [Schlobach & Cornet, 2003] Schlobach, S. and Cornet, R. (2003). Non-standard reasoning services for the debugging of description logic terminologies. In *International Joint Conference on Artificial Intelligence*, volume 18, pages 355–362. Citeseer.
- [Wiser & Carey, 1983] Wiser, M. and Carey, S. (1983). When heat and temperature were one. In Stevens, A. and Gentner, D., (eds.), *Mental Models*, pages 267–297. Erlbaum.