GALILEO: A System for Automating Ontology Evolution

GALILEO: A System for Automating Ontology Evolution

Michael Chan **Jos Lehmann** **Alan Bundy**
School of Informatics, University of Edinburgh
Edinburgh, U.K.
{mchan, jlehmann, bundy}@inf.ed.ac.uk

Abstract

We describe the GALILEO system, which is designed for automating the evolution of higher-order logic ontologies by incorporating user interaction for diagnosing and repairing faults. In particular, we present our approach to ontological conflict diagnosis, which circumvents problems posed by HOL's undecidability by means of: formalising modular ontologies as Isabelle *locales* and preparing the system by user interaction; applying *ontology repair plans* and automatically identifying logically valid terms that are responsible for the conflict; and, automatically eliminating physically meaningless terms.

1 Introduction

The process of revising an ontology in the face of new information is key to many areas of Computer Science. The literature on the subject calls such processes *ontology evolution* and concentrates on how Description Logic (DL) ontologies for Semantic Web applications need to evolve in order to maintain their coherence and consistency (ontology debugging) or relative to other ontologies (ontology alignment). We decided to investigate ontology evolution by *formalizing and automating in higher-order logic* (HOL) the mechanisms to repair *locally consistent* but possibly *globally inconsistent* ontologies, i.e. ontologies that are individually consistent but may give rise to an inconsistency when merged. Working with locally consistent ontologies enables the reasoning about the shape of the cause of the global conflict, allowing for specific meaningful repairs. The notion of ontologies we adopted is a general one, i.e. a specification of a conceptualisation [Gruber, 1993], which is compatible with our treatment of HOL theories as ontologies. Unlike the less expressive logics, including DL, our approach naturally allows for the formalisation of ontology evolution both as belief revision and syntactic manipulation (e.g., splitting a function, changing its arity, etc.). HOL has proven advantageous in at least three other ways. Firstly, HOL's polymorphism of variables (and other symbols such as $\leq$, $\geq$, $+$, $-$, etc.) permits the generality of the proposed repair plans for evolution and their applicability over diverse cases. Secondly, HOL-based theorem provers, such as Isabelle, enable HO-reasoning for ontology evolution and reasoning over locally consistent but globally inconsistent ontologies that share variables. Finally, many complex concepts are better represented as HOL objects, e.g., the orbit of a star – this is relevant because the examples used for developing and testing GALILEO are based on the evolution of physics, which involves concepts best represented as functions. This, therefore, relates our work to scientific discovery [Langley, 1981].

In this paper, we provide an overview of the GALILEO system, which is the first system to attempt to automate ontology evolution in HOL. In particular, we present our approach to ontological conflict diagnosis, which circumvents problems posed by HOL's undecidability by means of: a) formalising modular ontologies as *locales* [Ballarin, 2006] and preparing the system by user interaction; b) automatically identifying logically valid terms that are responsible for the conflict; and, c) automatically eliminating physically meaningless terms. We have designed mechanisms called *ontology repair plans* (ORPs) [Bundy and Chan, 2008], each of which defines a trigger pattern for detecting that certain faults exist in the given ontologies and a set of transformation rules for resolving the detected conflict. The identification of the term that gives rise to the inconsistency, hereafter designated by "conflict term" and the variable *stuff*, underpins the diagnosis component and is a relatively complex reasoning task – whereas the repair component typically involves relatively simple syntactic manipulation of the ontologies based on the result of the diagnosis. Thus, in this paper, we present only GALILEO's diagnosis component and briefly outline the repair.

2 Case Study and Ontology Repair Plans

We describe below the running example for the rest of the paper: the contradiction between the predicted and the observed orbital velocity of galaxies. We then describe two ORPs, which can be applied to resolve the conflict that underlies the case study and produce two different but historically proposed explanations.

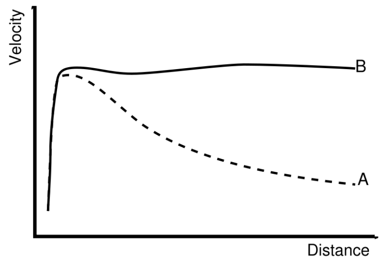

Figure 1: The x-axis is the distance of the stars in the galaxy $Glxy1$ from the galactic centre and the y-axis is their orbital velocities. The dotted and solid lines are the predicted and actual curve that is observed.[a]

[a] Taken from `http://en.wikipedia.org/wiki/Galaxy_rotation_problem`.

2.1 Case Study: Predicted vs. Observed Orbital Velocity of Galaxies

Newtonian dynamics predicts that orbital velocities[1] of stars in a spiral galaxy decrease inversely with the square root of the distance from the galactic centre, or the *radius*. However, the velocity observed was almost constant out to large radii (Figure 1), as first observed by Rubin [Rubin *et al.*, 1980]. One explanation is about the theoretical existence of *dark matter*, which is matter that is inferred to exist in the universe but not give off light, and is based on various sources of evidence, including the fact that the observed orbital velocities of stars in spiral galaxies exceeds the predicted values. Rubin concluded that some invisible matter is exerting a gravitational force on these stars, causing the unexpectedly high velocities.

Suppose the contradiction arises from observing the galaxy rotation curves Ga and Gb of type $Gly \mapsto Real \mapsto Real$, which are the curves based on prediction and observation, respectively, where Gly is the type of galaxy. Let the ontology of the predictive theory Thy contain definitions, including that of orbital velocity, and assertions about the value of various properties of different stars; let the sensory ontology, $Sens$, contain observations made on the *rotational velocity* of a galaxy, as reported in [Rubin *et al.*, 1980]. Figure 2 highlights some of the axioms available in the two ontologies, where Evt, Gly, Dst, and Str denote types for representing events, galaxies, distances, and stars, respectively; G denotes the universal gravitational constant; OV, the orbital velocity; g_r, the set of stars in the galaxy g up to the radius r; M, the mass of a body; Rad, the radius of a star from its galactic centre; RtV, the rotational velocity; RdV, the radial velocity; SyV, the velocity of the galactic system relative to the observer; Inc, the inclination of the galaxy; $\lambda Shift$, the shift in wave length; $Glxy1$, the galaxy being observed; $Star1$ and $Star2$, stars in the observed galaxy; and, $Obs6$, the observation event.

Note that the graph object in (1) is defined using equality on functions, whereas that in (4) is defined using the equality

[1] We assume that observed orbital velocity and rotational velocity for stars are the same.

$$
\begin{aligned}
&Ax(Thy) := \{\, \forall e{:}Evt.\, Ga(Glxy1) := OV(e, Glxy1) && (1)\\
&\quad \forall e{:}Evt,\ g{:}Gly,\ r{:}Dst.\ OV(e,g) := && (2)\\
&\qquad \sqrt{\frac{G \times \sum s \in g_r.M(e,s)}{r}},\ \ G := 6.67 \times 10^{-11}, &&\\
&\quad M(Obs6, Star1) := 5,\ M(Obs6, Star2) := 9\ \ldots \} && (3)\\
&Ax(Sens) := \{\, \forall s{:}Str. && (4)\\
&\qquad Gb(Glxy1, Rad(s, Glxy1)) := &&\\
&\qquad RtV(Obs6, s, Glxy1), &&\\
&\quad \forall e{:}Evt,\ s{:}Str,\ g{:}Gly.\ RtV(e,s,g) := && (5)\\
&\qquad \frac{RdV(e,s,g) - SyV(e,g)}{sin(Inc(e,g))}, &&\\
&\quad \lambda Shift(Obs6, Star1, Glxy1, \ldots) = 300, \quad \ldots \} && (6)
\end{aligned}
$$

Figure 2: Example axioms for the representation of the case study in §2.1.

on reals; this illustrates a benefit of the polymorphism in HOL over less expressive logics. Moreover, we assume both Thy and $Sens$ contain some knowledge about shapes of curves, e.g., the curve that has a positive gradient between zero and some point and a zero gradient thereafter. A modular setup permits the extension and sharing of such ontological knowledge.

Some symbols in $Sens$, such as RtV and RdV, are not in the language of Thy. To link together the seemingly disparate terms, an example axiom of the bridging meta-ontology, O_b, is:

$$
\begin{aligned}
&\forall v, r{:}Real.\ (Sens\ \vdash\ Rad(s,g) = r\ \wedge && (7)\\
&Sens\ \vdash\ RtV(e,s,g) = v) \longleftrightarrow Thy\ \vdash\ OV(e,g,r) = v &&
\end{aligned}
$$

which relates RtV in $Sens$ to the OV in Thy by Rad. Rubin's proposed resolution is the hypothetical existence of dark matter, which introduces a new form of matter. In contrast, the *MOdified Newtonian Dynamics* (MOND) [Milgrom, 1983] argues that the definition of G is flawed and that it is not a constant, but is a function depending on the acceleration of the star. By MOND, the value of G is greater for stars with very low accelerations.

2.2 Ontology Repair Plans

We introduce here two ORPs, *Where's My Stuff* and *Inconstancy*, and focus mostly on the diagnostic aspect, i.e. their trigger formulae.

The *Where's My Stuff* Ontology Repair Plan

In physics, a common type of conflict is caused by a difference between the predicted value of some property and the value according to some corresponding sensory information arising from an experiment. This conflict is typically caused by the use of a theoretical definition that is based on an incorrect definition of the property measured or that of a dependency of the property measured. The error is generally a consequence of a misconceptualisation with a component of the property neglected from the definition; for instance, when speaking about the matter distribution inside a galaxy, only the visible matter is included in the theoretical definition and dark matter is neglected.

Suppose an ontology O_1, representing the current state of a predictive theory, disagrees over the value of $f(\mathit{stuff})$ with another ontology O_2, representing some sensory information. $f(\mathit{stuff})$ might, for instance, be the graph plotting the orbital velocities of stars – a HO-object.

Trigger: If $f(\mathit{stuff})$ has different values, the following formula will be triggered:

$$\exists O_1, O_2{:}\mathit{Onto},\ \tau, \tau'{:}\mathit{Types},\ f{:}\tau \mapsto \tau', \tag{8}$$
$$\mathit{stuff}{:}\tau,\ v{:}\tau'.$$
$$(O_1 \vdash f(\mathit{stuff}) >_{\tau'} v \wedge O_2 \vdash f(\mathit{stuff}) \leq_{\tau'} v) \tag{9}$$
$$\vee$$
$$(O_1 \vdash f(\mathit{stuff}) \geq_{\tau'} v \wedge O_2 \vdash f(\mathit{stuff}) <_{\tau'} v) \tag{10}$$

where $O \vdash \phi$ means that ϕ is a theorem of ontology O; $t{:}Types$ means t is a type; $o{:}Onto$ means o is an ontology; v is a value of type τ'; and, $>_\tau$ is a partial order for τ.

WMS is triggered if the return value of $f(\mathit{stuff})$ deduced from O_1 is different from that deduced from O_2 (9, 10). The two variables in $f(\mathit{stuff})$, stuff and f, respectively represent the part of the term that is responsible and not responsible for the conflict. WMS resolves the detected conflict by splitting stuff into three parts: visible stuff, invisible stuff, and total stuff, and defining invisible stuff in terms of total and visible stuffs in the repaired O_1, $\nu(O_1)$. The new O_2, $\nu(O_2)$, is the same as O_2 except for the renaming of old stuff to visible stuff. If Thy and $Sens \cup O_b$ in §2.1 were instantiations of O_1 and O_2, respectively, and stuff was instantiated to $\mathit{Glxy1}$, where $\cup$ merges ontologies, then WMS emulates Rubin's solution by introducing a third component to the galaxy: dark matter.

2.3 The *Inconstancy* Ontology Repair Plan

In order to better define a repair operation, so that, e.g., the number of possible values a new variable created by the repair process can take is dramatically reduced, extra sensory information can help resolve such ambiguity. We have identified from historical records of physics development that the inclusion of additional sensory information collected under different conditions can indeed help identify new dependencies for a property.

Suppose an ontology O_1, representing the current state of a physical theory, is in conflict with a set of ontologies $O_2(\vec{r_i})$, representing sensory information arising from experiments under different circumstances, represented by $\vec{r_i}$, such that unexpected variations in the value of $f(\mathit{stuff})$ arise under these circumstances.

Trigger: If $f(\mathit{stuff})$ is measured to take different values in different circumstances, then the following trigger formulae will be matched:

$$\exists O_1, O_2{:}\mathit{Onto},\ \tau, \tau'{:}\mathit{Types},\ \vec{r_1}, \vec{r_2}{:}\vec{\mathit{Cir}},$$
$$f{:}\tau \mapsto \tau', \mathit{stuff}{:}\tau,\ \vec{s_1}, \vec{s_2}{:}\tau,\ c_1, c_2{:}\tau'.$$
$$O_1 \vdash \mathit{stuff} =_{\tau'} c\ \wedge \tag{11}$$
$$O_2(\vec{r_1}) \vdash f(\mathit{stuff}) =_{\tau'} c_1\ \wedge\ \ldots \tag{12}$$
$$O_2(\vec{r_n}) \vdash f(\mathit{stuff}) =_{\tau'} c_n\ \wedge$$
$$\exists i, j \leq n.\ O_1 \vdash c_i \neq_{\tau'} c_j\ \vee\ c \neq_{\tau'} c_i) \tag{13}$$

where $O_2(\vec{r})$ is the sensory ontology containing observations made under the condition $\vec{r}$; Cir is the type of conditions describing circumstances; $\vec{s_1}$ and $\vec{s_2}$ are arguments of stuff; and, c_1, c_2 are values of type τ'.

Inconstancy is triggered if $f(\mathit{stuff})$ is predicted to be independent of $\vec{r}$, but the return values of $f(\mathit{stuff})$ unexpectedly vary when $\vec{r}$ varies (11, 12, 13). The resolution of the detected conflict is to retain all O_1-axioms in the repaired O_1, $\nu(O_1)$, and retain all $O_2(\vec{r_i})$-axioms in the repaired $O_2(\vec{r_i})$, $\nu(O_2(\vec{r_i}))$, except for the replacement of old stuff with $\nu(\mathit{stuff})$ and the replacement of the definition of stuff by a new definition of $\nu(\mathit{stuff})$ in $\nu(O_1)$ – the new definition establishes a relationship between $\vec{r_i}$ and stuff. If G in §2.1 is the instantiation of stuff, then G is given a new dependency, which is the solution proposed by MOND.

3 The GALILEO System

Both WMS and Inconstancy have been implemented in the GALILEO system. The reasoning capabilities of GALILEO are provided by the Isabelle proof assistant [Paulson, 1994], so soundness of the reasoning is guaranteed. Beside reasoning, Isabelle provides a contextualisation facility called *locales* [Ballarin, 2006], which can be used for formalising modular ontologies. The GALILEO system handles flexible ontological configurations, including ontological extensions and those with heterogeneous signatures. For heterogeneous configurations, *bridging ontologies* are used to align signatures across different ontologies. GALILEO analyses ontologies expressed in an Isabelle theory file, in which each ontology is represented as a locale.

There are three major phases in the diagnosis procedure: a) *preparation*, in which the system acquires the instantiation of some of the existential variables in (9, 10) or (11, 12, 13) in order to detect a WMS type or an Inconstancy type of fault; b) *conflict term search*, in which all logically correct instantiations of stuff are discovered; and, c) *search space control*, in which only physically meaningful instantiations are retained. Below, we present our approach in more depth by applying it to an example record of ontology evolution in physics.

3.1 Preparation

In order to detect a WMS type of fault, the system requires an indication of the ontologies that may contain the fault, which represents the instantiations of O_1 and O_2. The system also requires a formula that expresses the conflict, $cexp$, which is a sentence that is derived to be true in one ontology but false in another. For example, in Rubin's graph of observed orbital velocities at various radii in a galaxy glx, the observed velocities near the edge of glx are greater than those predicted

[Rubin *et al.*, 1980]. Again, let Sens be the ontology containing sensory data and $Gb(glx)$ be the observed graph for glx. The diagnosis procedure begins with the substitution:

$$\{O_1/\mathit{Thy}, O_2/\mathit{Sens}, \mathit{cexp}/Ga(glx) = Gb(glx)\}$$

since $\mathit{Thy} \vdash Ga(glx) = Gb(glx)$ and $\mathit{Sens} \vdash Ga(glx) < Gb(glx)$ as Ga and Gb are expected to be identical but Gb turns out to be greater than Ga. To ensure that the two graphs indeed disagree, the user must discharge a proof obligation in the shape of the trigger formula but with the existentials instantiated, e.g.

$$(\mathit{Thy} \longrightarrow Ga(glx) \geq Gb(glx) \wedge \\ \mathit{Sens} \longrightarrow Ga(glx) < Gb(glx)) \vee \ldots$$

in which each ontology is treated as a predicate expressed as a conjunction of its axioms. Although the system requires values of O_1, O_2 and cexp as user inputs, the instantiations of each of the remaining variables, i.e. f, stuff, and v, are automatically identified, so the rest of the procedure is completely automatic.

For Inconstancy to be triggered, pattern (11, 12, 13) needs to be instantiated. Similar to WMS, GALILEO requires from the user the ontologies that may give rise to a fault – in this case, one theoretical ontology and two sensory ontologies. In addition, the system requires the terms representing the unexpected variation in each of the two input sensory ontologies, cexp_1 and cexp_2, respectively. For instance, suppose in two sensory ontologies $\mathit{Obs_1}$ and $\mathit{Obs_2}$, the observed orbital velocities of two stars, p_1 and p_2, unexpectedly vary. So, $\mathit{Obs_1} \vdash OV(p_1, glx) = v_1$ and $\mathit{Obs_2} \vdash OV(p_2, glx) = v_2$, and in the theory, $\mathit{Thy} \vdash OV(p_1, glx) - v_1 \neq OV(p_2, glx) - v_2$. Thus, we have the substitution:

$$\{O_1/\mathit{Thy}, O_1(\vec{r_1})/\mathit{Obs_1}, O_2(\vec{r_2})/\mathit{Obs_2}, \\ \mathit{cexp}_1/OV(p_1, glx), \mathit{cexp}_2/OV(p_2, glx)\}.$$

Like WMS, the fault is verified by discharging a proof obligation in the shape of the trigger formula, and the system then identifies various instantiations of stuff.

3.2 Search for Conflict Terms

Having indicated the potentially faulty ontologies and cexp, GALILEO identifies instantiations of stuff to completely instantiate the relevant trigger pattern. Isabelle's polymorphism is partial as its type unification does not specialise schematic type variables to function types. Such a constraint, which also holds in most other theorem provers, renders type unification tractable, but it is an impediment to instantiating stuff to different arity functions in different case studies. Letting stuff be a variadic function allows for high generality of the trigger. GALILEO, therefore, explicitly generates a set of types for stuff, given a maximum function order and arity.

With WMS, instantiations of f, stuff, and v can be discovered by applying Isabelle's higher-order matching algorithm to match cexp against the pattern $f(\mathit{stuff}) = v$. For example, there exists the substitution $\{glx/\mathit{stuff}\}$ by matching with $Ga(glx)$, which suggests that the conceptualisation of the galaxy glx should be adjusted, e.g., by introducing dark matter. Unlike the HOL unification problem, HOL matching is decidable.

Inconstancy also adopts a general approach to the identification of the conflict term by matching cexp_1 and cexp_2 against the patterns $f(\mathit{stuff}_1)$ and $f(\mathit{stuff}_2)$. The system iterates through each substitution returned and generates a proof obligation for each instantiation of stuff. With the substitution for Inconstancy described in §3.1, instantiations of stuff_i could be simply $OV(p_i, glx)$, then the proof obligation $\exists c_1, c_2.\ \mathit{Obs_1} \vdash OV(p_1, glx) = c_1 \wedge \mathit{Obs_2} \vdash OV(p_2, glx) = c_2 \wedge \mathit{Thy} \vdash c_1 \neq c_2$ will need to be discharged.

Often, the term expressed in cexp_i may not represent the preferred concept for repair, e.g., cexp_i represents measurements on orbital velocity but the preferred concept for repair is another, such as the gravitational constant G. To search for other potential concepts for repair, GALILEO automatically identifies other relevant instantiations of stuff. For example, Thy contains a definition of OV in terms of G and M (2); if stuff is instantiated to OV, then stuff can be instantiated to G, which is the fix proposed by MOND.

3.3 Search Space Control

For both ORPs, the search spaces of possible substitutions are typically unmanageable because f and stuff are both polymorphic – there are more than 20,000 substitutions in both cases, but only a small fraction of these are physically meaningful. A solution is needed to avoid an explosion of substitutions.

Isabelle's implementation of higher-order matching, like all others, returns logically correct substitutions without assessing the physical meaningfulness or relevance of the substitutions. Physical meaningfulness is crucial for controlling the search for suitable substitutions in order to perform a meaningful diagnosis, which must be compatible with real-world semantics. For instance, matcher 1 in Table 1 assigns stuff to $\mathit{Glxy1}$, which is the diagnosis required. However, we argue that some other example matches have little physical content or meaningfulness, e.g., matches 6 and 7 assign stuff to a function containing the identity function and a function containing a free variable ($?s$) generated by the matching algorithm, respectively.

1.	$\mathit{Glxy1}$	5.	$\lambda x.x(Ga(\mathit{Glxy1}))$
2.	Ga	6.	$\lambda x.x(Ga(x(\lambda x.\ x)))$
3.	$Ga(\mathit{Glxy1})$	7.	$\lambda x.x(Ga(x(?s(x))))$
4.	$\lambda x.\ x(\mathit{Glxy1})$		

Table 1: Example instantiations of stuff from pattern matching $f(\mathit{stuff})$ against $Ga(\mathit{Glxy1})$.

To reduce the number of substitutions returned, several heuristics (Table 2) are designed and applied.

4 Evaluation

We evaluated our technique and implementation on a diverse range of records of ontology evolution in physics. Although only one example and two different fixes have been described in this paper, we base the evaluation also on examples already reported elsewhere, e.g., the *bouncing-ball paradox* [Chan *et al.*, 2010] and the *discovery of latent heat* [Bundy *et al.*,

#	Description
H1	$stuff$ must not contain the $\lambda x.x$ term.
H2	f and $stuff$ must not contain free variables.
H3	$stuff$ must contain one or more constants.
H4	Two functions of the same function symbol but with permuted argument types are the same.

Table 2: Example heuristics for pruning the solution space: H1 ignores instantiations containing the identity function, which has no physical characteristic; H2 ignores instantiations that imply less ontological commitment than that we intended; H3 requires that the term to be repaired must contain an element of the signature, i.e. a concrete concept; and, H4 regards, e.g., $stuff : \alpha \to \beta \to \gamma$ and $stuff : \beta \to \alpha \to \gamma$ to be the same.

2010]; the goal is to demonstrate the high level of generality of our approach, not just being able to identify various types of $stuff$, but also to handle examples across distant subfields of physics. The rest of this section covers two key areas for evaluating our work: effectiveness of search space control and meaningfulness of the generated solutions.

4.1 Reduction in Candidate Matches

Our experiment has shown that our approach to automating conflict term discovery is intractable without a reduction in the search space containing candidate matches. Even with a low unification bound, e.g., 3, the procedure fails to return a solution on all case studies. The raw search spaces are, therefore, far too vast in practice. The use of the few example heuristics in Table 2, however, prunes away from the solution scope more than 20,000 matches for the case study in §2.1 and 57,000 for both the bouncing-ball and the latent heat case studies, and enables termination. The resulting sets of matches are significantly reduced to substantially more manageable amounts of 5, 21, and 27 for each respective example. Thus, the results suggest that the heuristics used are effective in controlling the search.

4.2 Meaningfulness of Solutions

Even though the solution spaces could be considerably reduced, we will examine the degree of physical meaningfulness of the instantiations of $stuff$ in the remaining matches. For dark matter, only matches 1 to 5 in Table 1 remain in the pruned space. Every candidate match returned is in fact physically interpretable – for instance, instantiating $stuff$ to $Glxy1$ indicates that the galaxy should have an extra component; to Ga leads to a redefinition of all predicted curves; and, to $Ga(Glxy1)$ limits the modification to only the definition of the predicted curve for $Glxy1$.

5 Conclusion

The paper provided an overview of the GALILEO system and of two of its ORPs: WMS and Inconstancy. The diagnostic components were presented by successfully applying them to discover two fault types underlying the same set of conflicting observations in astronomy. This showed the generality of GALILEO and its ability to return only physically meaningful ontology repairs.

References

[Ballarin, 2006] Clemens Ballarin. Interpretation of locales in isabelle: Theories and proof contexts. In *Mathematical Knowledge Management*, volume LNCS 4108, pages 31–43. Springer Berlin/Heidelberg, 2006.

[Bundy and Chan, 2008] Alan Bundy and Michael Chan. Towards ontology evolution in physics. In Wilfrid Hodges and Ruy de Queiroz, editors, *Logic, Language, Information and Computation*, volume 5110 of *Lecture Notes in Computer Science*, pages 98–110. Springer Berlin / Heidelberg, July 2008.

[Bundy *et al.*, 2010] A. Bundy, J. Lehmann, G. Qi, and I. Varzinczak, editors. *Notes of the ECAI-10 Workshop on Automated Reasoning about Context and Ontology Evolution*. ECAI-10, 2010.

[Chan *et al.*, 2010] M. Chan, J. Lehmann, and A. Bundy. Higher-order representation and reasoning for automated ontology evolution. In *Proceedings of the 2010 International Conference on Knowledge Engineering and Ontology Development*, pages 84–93. SciTePress, 2010.

[Gruber, 1993] Thomas R. Gruber. A translation approach to portable ontology specification. *Knowledge Acquisition*, 5(2), 1993.

[Langley, 1981] P. Langley. Data-driven discovery of physical laws. *Cognitive Science*, 5(1):31–54, 1981.

[Milgrom, 1983] M. Milgrom. A modification of the Newtonian dynamics as a possible alternative to the hidden mass hypothesis. *The Astrophysical Journal*, 270:365–370, 1983.

[Paulson, 1994] L. C. Paulson. *Isabelle: A generic theorem prover*. Springer-Verlag, 1994.

[Rubin *et al.*, 1980] V. C. Rubin, N. Thonnard, and Jr. Ford, W. K. 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, 1980.