

First-Order Default Logic Revisited

Yi Zhou

Artificial Intelligence Research Group
School of Computing, Engineering and Mathematics
University of Western Sydney, Australia

Abstract

Reiter's original proposal for default logic is unsatisfactory for open default theories because of Skolemization and grounding. In this paper, we reconsider this long-standing problem and propose a new world view semantics for first-order default logic. Roughly speaking, a world view of a first-order default theory is a maximal collection of structures satisfying the default theory where the default part is fixed by the world view itself. We show how this semantics generalizes classical first-order logic and first-order answer set programming, and we discuss its connections to Reiter's semantics and other related semantics. We also argue that first-order default logic under the world view semantics provides a rich framework for integrating classical logic based and rule based formalisms in the first-order case.

Introduction

As a predominant approach for nonmonotonic reasoning, default logic has attracted many researchers since Reiter's seminal work (Reiter 1980). Reiter's semantics targets on general first-order default theories in two steps. First, an extension semantics is defined for closed default theories in which no free variable occurs in any default rules. Then, general default theories with free variables (also called open default theories) are converted into closed ones by Skolemization and grounding so that the above extension semantics can be applied.

However, as pointed out by Poole (1987) and many other researchers (Lifschitz 1990; Baader and Hollunder 1995; Kaminski 1995; 1999), Reiter's semantics may lead to counterintuitive phenomena for open default theories. For instance, Lifschitz (1990) proposed an example with a single default rule $\frac{M \neg P(x)}{\neg P(x)}$ and a single first-order assertion $P(a)$. One should expect that we have the following information hidden in the default theory

$$\forall x (P(x) \leftrightarrow (x = a)).$$

However, in Reiter's extension semantics, the default rule will only be grounded on a and another Skolem constant c different than a . Therefore, one can only conclude $P(a) \wedge \exists x \neg P(x)$.

Copyright © 2014, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

The unsatisfactory performance of Reiter's original semantics for open first-order default theories is indeed a serious semantic issue. As pointed out by Reiter himself (1980),

“the genuinely interesting cases involve open defaults.”

In this paper, we reconsider this long standing problem and propose a new world view semantics. Instead of first-order theories, we use collections of structures sharing the same domain and function interpretations (called *views*) as the candidate solution concept. Roughly speaking, a world view of a first-order default theory is a maximal view (collection of structures) satisfying the default theory where the default part is fixed by the world view itself. Although related, this semantics is essentially different from Reiter's extension semantics, even for closed default theories. However, they coincide on some restricted classes.

Our work is also strongly motivated from the attempt to combine classical logic based formalisms and rule based formalisms, particularly the recent development in ontology engineering for adding rules onto the description logics layer (Baader and Hollunder 1995; Motik et al. 2006; Eiter et al. 2008; Motik and Rosati 2010; Lukasiewicz 2010; Zhou and Zhang 2012). We argue that default logic provides a rich framework for this task. First, we show that it generalizes both classical logic and first-order answer set programming. Then, we show how default logic can handle both monotonic reasoning and nonmonotonic reasoning and flexibly switch between open world reasoning and closed world reasoning.

The main contribution of this paper is to provide a new world view semantics for first-order default logic. Although this is a long standing problem, we believe that research in this direction will further deepen our understandings and shed new insights in

- default logic and nonmonotonic reasoning,
- dealing with free variables in complex knowledge representation formalisms, and
- integration of classical logic based and rule based formalisms in the first-order case.

Syntax

We assume that the readers are familiar with the basic notions and notations in classical first-order logic. A (first-

order) *default rule* (rule for short if clear from the context) is of the following form

$$\alpha_1 \mid \dots \mid \alpha_n \leftarrow \beta_1, \dots, \beta_m, \text{Not } \gamma_1, \dots, \text{Not } \gamma_l, \quad (1)$$

where $\alpha_i, 1 \leq i \leq n$, $\beta_j, 1 \leq j \leq m$ and $\gamma_k, 1 \leq k \leq l$ are first-order formulas. A *default theory* is a set of rules of the form (1).

Intuitively, rule (1) means that if an agent believes that every $\beta_j, 1 \leq j \leq m$ is true, and the agent does not believe that any of $\gamma_k, 1 \leq k \leq l$ is true, then the agent must believe that at least one of $\alpha_i, 1 \leq i \leq n$ is true.

Given a rule r of the form (1), the set $\{\alpha_1, \dots, \alpha_n\}$ is called the *head* of r , denoted by $\text{Head}(r)$; the set $\{\beta_1, \dots, \beta_m\}$ is called the *positive body* of r , denoted by $\text{Pos}(r)$; the set $\{\text{Not } \gamma_1, \dots, \text{Not } \gamma_l\}$ is called the *negative body* of r , denoted by $\text{Neg}(r)$ ¹; also the set $\{\beta_1, \dots, \beta_m, \text{Not } \gamma_1, \dots, \text{Not } \gamma_l\}$ is called the *body* of r , denoted by $\text{Body}(r)$. Clearly, $\text{Body}(r) = \text{Pos}(r) \cup \text{Neg}(r)$.

Notice that the form (1) follows the well adopted syntactic form of rules in logic programming (in particular answer set programming), which actually differs from Reiter's original syntax of default rules as follows

$$\frac{\alpha(\vec{x}) : M\beta_1(\vec{x}), \dots, M\beta_m(\vec{x})}{\omega(\vec{x})}, \quad (2)$$

where $\alpha(\vec{x}), \beta_i(\vec{x}), 1 \leq i \leq m$ and $\omega(\vec{x})$ are formulas and \vec{x} is the set of free variables occurred in all these formulas. Intuitively, a rule of form (2) means that if an agent believes that $\alpha(\vec{x})$ is true and her/his beliefs are consistent with any of $\beta_j, 1 \leq j \leq m$, then the agent must believe that $\omega(\vec{x})$ is true. In addition, in Reiter's original syntax, a default theory is a pair (D, W) , where D is a set of rules of the form (2) and $W = \{S_1, \dots, S_k, \dots\}$ is a set of sentences.

Our syntax is actually a generalization of Reiter's syntactic form in the sense that a rule of the Reiter's form

$$\frac{\alpha(\vec{x}) : M\beta_1(\vec{x}), \dots, M\beta_m(\vec{x})}{\omega(\vec{x})}, \quad (3)$$

can be rewritten as

$$\omega(\vec{x}) \leftarrow \alpha(\vec{x}), \text{Not } \neg\beta_1(\vec{x}), \dots, \text{Not } \neg\beta_m(\vec{x}) \quad (4)$$

in form (1), and a sentence $S \in W$ in a default theory (D, W) can be simply rewritten as

$$S \leftarrow$$

in form (1). Notice that the component $M\beta_j(\vec{x})$ in form (2) is actually corresponding to $\text{Not } \neg\beta_j(\vec{x})$ but not $\text{Not } \beta_j(\vec{x})$ in (4). Here, \neg is the classical negation operator.

We say that a rule r of the form (1) is a *fact* if $\text{Pos}(r) = \text{Neg}(r) = \emptyset$, a *constraint* if $\text{Head}(r) = \emptyset$, *positive* if $\text{Neg}(r) = \emptyset$, a *logic programming rule* if all formulas occurred in r (i.e. $\alpha_i, 1 \leq i \leq n$, $\beta_j, 1 \leq j \leq m$, $\gamma_k, 1 \leq k \leq l$) are atoms, *closed* if all formulas occurred in r are closed.

We say that a default theory is a *classical theory* if every rule in it is a fact, a *logic program* if every rule in it is a logic programming rule, *normal* if every rule in it is normal, *closed* if every rule in it is closed.

¹The head, the positive body and the negative body of a default rule are also called consequent/conclusion, prerequisite and justification respectively in the literature.

Defining the World View Semantics

We start to define the world view semantics for first-order default logic.

Definition 1 (Configuration). Let τ be a vocabulary. A *configuration* of τ is a tuple $\Delta = (A, c_1^A, \dots, c_m^A, f_1^A, \dots, f_l^A)$, where A is the domain, c_1^A, \dots, c_m^A are the constant mappings and f_1^A, \dots, f_l^A are the function interpretations.

A configuration Δ can be extended into a structure \mathcal{M} by interpreting the predicates. In this case, we say that Δ is the configuration of \mathcal{M} , and conversely, \mathcal{M} is a structure on Δ . We use $\Delta^{\mathcal{M}}$ to denote the configuration of \mathcal{M} . Then notion of configuration is analogous to the notion of pre-interpretation introduced by Kaminski (1995).

Definition 2 (View). A *view* of a vocabulary τ is a collection of τ -structures with the same configuration, i.e., sharing the same domain, constant mappings and function interpretations. More precisely, let Δ be a configuration, a view \mathcal{W} on Δ is a collection $\mathcal{M}_1, \dots, \mathcal{M}_t, \dots$ of structures such that for all t , $\Delta^{\mathcal{M}_t} = \Delta$.

Here, we say that Δ is the configuration of \mathcal{W} . Sometimes we use $\Delta^{\mathcal{W}}$ to denote the configuration of \mathcal{W} for better readability. Intuitively, a view represents an agent's belief in the sense that each structure in the view is corresponding to a possible world.

We now define the satisfaction relation between views associated with assignments and default rules. Let r be a rule of the form (1), $\mathcal{W} = \{\mathcal{M}_1, \dots, \mathcal{M}_t, \dots\}$ a view and η an assignment. We say that \mathcal{W} *satisfies* r on η , denoted by $\mathcal{W} \models r[\eta]$, if

$$(\dagger) \mathcal{W} \models \text{Pos}(r)[\eta] \text{ and } \mathcal{W} \not\models \text{Neg}(r)[\eta] \text{ implies that } \mathcal{W} \models \text{Head}(r)[\eta].$$

As the default part (i.e. the negative body) of a default rule plays a role of assumption, we also need a different notion to define the satisfaction relation by fixing this part. Let r be a rule of the form (1), $\mathcal{W} = \{\mathcal{M}_1, \dots, \mathcal{M}_t, \dots\}$ and $\mathcal{W}' = \{\mathcal{M}'_1, \dots, \mathcal{M}'_t, \dots\}$ two views, and η an assignment. We say that \mathcal{W}' *satisfies* r on η by fixing the default part with \mathcal{W} , denoted by $(\mathcal{W}', \mathcal{W}) \models r[\eta]$, if

$$(\ddagger) \mathcal{W}' \models \text{Pos}(r)[\eta] \text{ and } \mathcal{W}' \not\models \text{Neg}(r)[\eta] \text{ implies that } \mathcal{W}' \models \text{Head}(r)[\eta].$$

Let Π be a default theory, η an assignment and \mathcal{W} and \mathcal{W}' two views. We say that \mathcal{W} *satisfies* Π on η , denoted by $\mathcal{W} \models \Pi[\eta]$, if for all $r \in \Pi$, $\mathcal{W} \models r[\eta]$. We say that \mathcal{W} *satisfies* Π , denoted by $\mathcal{W} \models \Pi$, if for all assignments η , $\mathcal{W} \models \Pi[\eta]$. We say that \mathcal{W}' *satisfies* Π on η by fixing the default part with \mathcal{W} , denoted by $(\mathcal{W}', \mathcal{W}) \models \Pi[\eta]$, if for all $r \in \Pi$, $(\mathcal{W}', \mathcal{W}) \models r[\eta]$. We say that \mathcal{W}' *satisfies* Π by fixing the default part with \mathcal{W} , denoted by $(\mathcal{W}', \mathcal{W}) \models \Pi$, if for all assignments η , $(\mathcal{W}', \mathcal{W}) \models \Pi[\eta]$.

The following property holds straightforwardly from the definitions.

Proposition 1. Let Π be a default theory and \mathcal{W} a view. Then, $(\mathcal{W}, \mathcal{W}) \models \Pi$ iff $\mathcal{W} \models \Pi$.

Definition 3 (World View). Let Π be a default theory and \mathcal{W} a view. We say that \mathcal{W} is a *world view* of Π if $\mathcal{W} \models \Pi$

and there does not exist another view \mathcal{W}' such that $\mathcal{W} \subset \mathcal{W}'$ and $(\mathcal{W}', \mathcal{W}) \models \Pi$.

Intuitively, a world view of a default theory is a maximal collection of structures satisfying the default theory, where the default part is fixed by the world view itself. Definition 3 can be presented in an alternative way closer to Reiter's original fixed-point definition. Let Π be a default theory and \mathcal{W} a view such that $\mathcal{W} \models \Pi$, by $\Gamma(\mathcal{W})$, we denote the maximal view (in terms of set inclusion) satisfying the following conditions:

- $\mathcal{W} \subseteq \Gamma(\mathcal{W})$.
- $(\Gamma(\mathcal{W}), \mathcal{W}) \models \Pi$.

Notice that a drawback of this definition is that one needs to assure that such a $\Gamma(\mathcal{W})$ always exists (although this is indeed the case under our context). Clearly, \mathcal{W} is a world view of Π iff \mathcal{W} is the fixed point of the operator Γ , i.e., $\mathcal{W} = \Gamma(\mathcal{W})$.

The world view semantics has many origins from the literature. The term world view is borrowed from (Gelfond 1994), although the semantics is defined in a very different manner. Definition 3 actually follows the well known Gelfond-Lifschitz reduct semantics for propositional answer set programming, that is, an answer set of a propositional logic program is a maximal set of propositional atoms satisfying the reduct of the program relative to this set. Here, $(\mathcal{W}', \mathcal{W}) \models \Pi$ is analogous to $\mathcal{W}' \models \Pi^{\mathcal{W}}$, where $\Pi^{\mathcal{W}}$ is the reduct of Π relative to \mathcal{W} . However, we cannot use the notation of reduct $\Pi^{\mathcal{W}}$ here in our world view semantics as it is not a real program or default theory.

Of course, Definition 3 shares some basic ideas of Reiter's original semantics. On one side, the default part plays a role of assumptions, which is fixed by a candidate. On the other side, the world view semantics can also be defined in a fixed point style as shown before. However, the main difference between the world view semantics and Reiter's original semantics is that the solution concept for the former is defined upon views (i.e. collections of structures), while the one for the latter is defined upon first-order theories (i.e. collections of sentences). Note that collections of structures and first-order theories correspond to each other in the sense that all the models of a first-order theory form a collection of structures. We shall discuss later in the paper why we need to define upon collections of structures rather than collections of sentences for first-order open default theories, and the detailed relationships between the world view semantics and Reiter's original one.

Relationships to Reiter's Semantics

In this section, we study the relationships between the world view semantics and Reiter's extension semantics (Reiter 1980). Clearly, for open default theories, the world view semantics is very different from Reiter's extension semantics. In fact, this is indeed the case even for closed default theories as well. The key reason why these two semantics behave quite differently, even for closed default theories, is that all views are forced to have the same configuration. This is because a configuration itself includes some information, for instance, the domain size and the mapping of terms.

Nevertheless, these two semantics coincide if we only consider the first-order sentences whose satisfiability is independent of a configuration. We say that a first-order sentence ϕ is *configuration independent* if that ϕ is consistent implies that ϕ has a model on every configuration. We say that a closed default theory is *configuration independent* if all formulas occurring in every rule in the default theory are configuration independent.

Theorem 1. *Let Π be a configuration independent closed default theory. If a first-order theory E is an extension of Π , then for any configuration Δ , the set of models of E on Δ is a world view of Π . Conversely, if a view \mathcal{W} on a configuration Δ is a world view of Π , then there exists an extension E of Π such that \mathcal{W} is the set of models of E on Δ .*

Corollary 2. *Let Π be a propositional default theory. Then, a formula F is equivalent to an extension of Π iff the set of models of F is a world view of Π .*

Classical First-Order Logic + First-Order Answer Set Programming

In this section, we show how the world view semantics generalizes both classical first-order logic and first-order answer set programming. We also argue that first-order default logic provides a rich framework for integrating them.

Theorem 3. *Let S be a set of first-order sentences and Δ a configuration such that S has models on Δ . Then, S has a unique world view on Δ , which is the class of all models of S on Δ .*

Next, we show that the stable model semantics for first-order answer set programming is a special case of the world view semantics for first-order default theories as well. A first-order logic program is a special case of first-order default rules except that the formulas occurring in rules are atoms. The semantics for first-order logic programs can be defined through by a transformation into second-order logic (Ferraris, Lee, and Lifschitz 2011; Zhou and Zhang 2011).

Theorem 4. *Let Π be a logic program. If a structure \mathcal{M} is a stable model of Π , then \mathcal{M}^+ is a world view of Π , where \mathcal{M}^+ contains all structures on the same configuration bigger than \mathcal{M} .² Conversely, if a view \mathcal{W} is a world view of Π , then the structure on the same configuration that interprets each predicate P as $\bigcap_{\mathcal{M} \in \mathcal{W}} P^{\mathcal{M}}$ is a stable model of Π .*

Theorems 3 and 4 show that both classical first-order logic and first-order answer set programming can be considered as special cases of first-order default logic. Moreover, we argue that default logic not only generalizes classical logic and answer set programming but also integrates them in a natural way. In the propositional case, propositional default logic provides an ideal framework for integrating classical propositional logic and propositional answer set programming. Lifted into the first-order case, first, the syntax of first-order default theory combines the syntax of both classical logic and logic programming. For semantics, Theorem 3 and

²We say that a structure \mathcal{M}' is bigger than another structure \mathcal{M} if for all predicates P , $P^{\mathcal{M}} \subseteq P^{\mathcal{M}'}$.

Theorem 4 show respectively that classical first-order logic and first-order disjunctive logic programming under the stable model semantics can be viewed as special cases of first-order default logic. Moreover, classical logic and logic programming actually form two different dimensions of first-order default logic, where the classical logic dimension is corresponding to the interrelationships among structures in a world view and the logic programming dimension is corresponding to the interrelationships among different world views.

Related Work

A number of alternative semantics have been proposed to address the unsatisfactory semantics for open defaults (Poole 1987; Lifschitz 1990; Baader and Hollunder 1995; Kaminski 1995; 1999). As far as we know, this problem was firstly reported by Poole (1987), who also provided an alternative solution by using Hilbert's ε -symbol. However, Baader and Hollunder (1995) pointed out that, although improved, this semantics is still not that satisfactory as it cannot escape from Skolemization. Instead, Baader and Hollunder considered a weak semantics, in which default rules only apply to known individuals.

The idea of using collections of structures as the solution concept for default logic is not new. This was presented by Guerreiro and Casanova for closed default theories (Guerreiro and Casanova 1990), and later extended by Lifschitz for open default theories restricted to a fixed universe (Lifschitz 1990). Lifschitz fixed a pre-given universe and extended the language by introducing a new constant for each element in the universe. An extension of a default theory is then defined as the set of sentences satisfied by a class of structures, which is a fixed point of an operator similar to Reiter's Γ operator but defined on classes of structures. Kaminski pointed out some problems with Lifschitz's approach, for instance, a default theory is not equivalent to its grounded version under the domain closure assumption (Kaminski 1995). Instead, Kaminski modified Lifschitz's semantics by fixing a pre-given infinite Herbrand universe as well as the interpretations of functions (Kaminski 1995).

In fact, the world view semantics can be regarded as a modification of Lifschitz's and Kaminski's approaches. However, there are three major differences. First, we do not introduce any new universe and we do not need to extend the language by introducing a new constant for each domain element in the new universe. Second, we do not ground on Herbrand universe nor on any other domains. Finally, we do not use the notion $Th(W)$ to denote the set of sentences satisfied by a structure or a class of structures. Although this notion is well defined, one has to be very careful as not every class of structures is corresponding to a set of first-order sentences. We regard the world view semantics as a purely model theoretical approach. Similar to the classical first-order semantics, the notions and notations used in the world view semantics are merely restricted in the language of the default theory.

Conclusion

The main contribution of this paper is the world view semantics for first-order default logic. Roughly speaking, a world view of a default theory is a maximal class of structures sharing the same configuration and satisfying the default theory, where the default part is fixed by the world view itself. A major rationale for using collections of structures as the solution concept is that first-order theories, such as extensions, are not powerful enough to capture all information hidden in a first-order default theory. The world view semantics is essentially different from Reiter's original proposal, even for closed default theories. Nevertheless, they coincide on configuration independent closed default theories. First-order default logic under the world view semantics provides a rich framework for knowledge representation and reasoning as it not only generalizes but also integrates classical first-order logic and first-order answer set programming.

References

- Baader, F., and Hollunder, B. 1995. Embedding defaults into terminological knowledge representation formalisms. *J. Autom. Reasoning* 14(1):149–180.
- Eiter, T.; Ianni, G.; Lukasiewicz, T.; Schindlauer, R.; and Tompits, H. 2008. Combining answer set programming with description logics for the semantic web. *Artificial Intelligence* 172(12-13):1495–1539.
- Ferraris, P.; Lee, J.; and Lifschitz, V. 2011. Stable models and circumscription. *Artificial Intelligence* 175(1):236–263.
- Gelfond, M. 1994. Logic programming and reasoning with incomplete information. *Ann. Math. Artif. Intell.* 12(1-2):89–116.
- Guerreiro, R., and Casanova, M. 1990. An alternative semantics for default logic. In *NMR '90*.
- Kaminski, M. 1995. A comparative study of open default theories. *Artif. Intell.* 77(2):285–319.
- Kaminski, M. 1999. Open default theories over closed domains. *Logic Journal of the IGPL* 7(5):577–589.
- Lifschitz, V. 1990. On open defaults. In *Computational Logic: Symposium Proceedings*, 230–237.
- Lukasiewicz, T. 2010. A novel combination of answer set programming with description logics for the semantic web. *IEEE Trans. Knowl. Data Eng.* 22(11):1577–1592.
- Motik, B., and Rosati, R. 2010. Reconciling description logics and rules. *Journal of the ACM* 57(5).
- Motik, B.; Horrocks, I.; Rosati, R.; and Sattler, U. 2006. Can OWL and logic programming live together happily ever after? In *International Semantic Web Conference*, 501–514.
- Poole, D. 1987. Variables in hypotheses. In *IJCAI*, 905–908.
- Reiter, R. 1980. A logic for default reasoning. *Artif. Intell.* 13(1-2):81–132.
- Zhou, Y., and Zhang, Y. 2011. Progression semantics for disjunctive logic programs. In *AAAI*.
- Zhou, Y., and Zhang, Y. 2012. RDL: Enhancing description logic with rules. In *AI*, 567–578.