

On Strong and Default Negation in Logic Program Updates

Martin Slota
CENTRIA
New University of Lisbon

Martin Baláz
Faculty of Mathematics, Physics and Informatics
Comenius University

João Leite
CENTRIA
New University of Lisbon

Abstract

Existing semantics for answer-set program updates fall into two categories: either they consider only *strong negation* in heads of rules, or they primarily rely on *default negation* in heads of rules and optionally provide support for strong negation by means of a syntactic transformation.

In this paper we pinpoint the limitations of both these approaches and argue that both types of negation should be first-class citizens in the context of updates. We identify principles that plausibly constrain their interaction but are not simultaneously satisfied by any existing rule update semantics. Then we extend one of the most advanced semantics with direct support for strong negation and show that it satisfies the outlined principles as well as a variety of other desirable properties.

1 Introduction

The increasingly common use of rule-based knowledge representation languages in highly dynamic and information-rich contexts, such as the Semantic Web (Berners-Lee, Hendler, and Lassila 2001), requires standardised support for updates of knowledge represented by rules. Answer-set programming (Gelfond and Lifschitz 1988; 1991) forms the natural basis for investigation of rule updates, and various approaches to answer-set program updates have been explored throughout the last 15 years (Leite and Pereira 1998; Alferes et al. 1998; 2000; Eiter et al. 2002; Leite 2003; Sakama and Inoue 2003; Alferes et al. 2005; Banti et al. 2005; Zhang 2006; Šeřfránek 2006; Delgrande, Schaub, and Tompits 2007; Osorio and Cuevas 2007; Šeřfránek 2011; Krümpelmann 2012).

The most straightforward kind of conflict arising between an original rule and its update occurs when the original conclusion logically contradicts the newer one. Though the technical realisation and final result may differ significantly, depending on the particular rule update semantics, this kind of conflict is resolved by letting the newer rule prevail over the older one. Actually, under most semantics, this is also the *only* type of conflict that is subject to automatic resolution (Leite and Pereira 1998; Alferes et al. 2000; Eiter et al. 2002; Alferes et al. 2005; Banti et al. 2005; Delgrande, Schaub, and Tompits 2007; Osorio and Cuevas 2007).

From this perspective, allowing for both *strong* and *default negation* to appear in heads of rules is essential for an expressive and universal rule update framework (Leite 2003). While strong negation is the natural candidate here, used to express that an atom *becomes explicitly false*, default negation allows for more fine-grained control: the atom only *ceases to be true*, but its truth value may not be known after the update. The latter also makes it possible to move between any pair of epistemic states by means of updates, as illustrated in the following example:

Example 1.1 (Railway crossing (Leite 2003)). *Suppose that we use the following logic program to choose an action at a railway crossing:*

cross \leftarrow \neg train. wait \leftarrow train. listen \leftarrow \sim train, $\sim\neg$ train.

The intuitive meaning of these rules is as follows: one should cross if there is evidence that no train is approaching; wait if there is evidence that a train is approaching; listen if there is no such evidence.

Consider a situation where a train is approaching, represented by the fact (train.). After this train has passed by, we want to update our knowledge to an epistemic state where we lack evidence with regard to the approach of a train. If this was accomplished by updating with the fact (\neg train.), we would cross the tracks at the subsequent state, risking being killed by another train that was approaching. Therefore, we need to express an update stating that all past evidence for an atom is to be removed, which can be accomplished by allowing default negation in heads of rules. In this scenario, the intended update can be expressed by the fact (\sim train.).

With regard to the support of negation in rule heads, existing rule update semantics fall into two categories: those that only allow for strong negation, and those that primarily consider default negation. As illustrated above, the former are unsatisfactory as they render many belief states unreachable by updates. As for the latter, they optionally provide support for strong negation by means of a syntactic transformation.

Two such transformations are known from the literature, both of them based on the principle of coherence: if an atom p is true, its strong negation $\neg p$ cannot be true simultaneously, so $\sim\neg p$ must be true, and also vice versa, if $\neg p$ is true, then so is $\sim p$. The first transformation, introduced in (Alferes and Pereira 1996), encodes this principle directly

by adding, to both the original program and its update, the following two rules for every atom p :

$$\sim\neg p \leftarrow p. \quad \sim p \leftarrow \neg p.$$

This way, every conflict between an atom p and its strong negation $\neg p$ directly translates into two conflicts between the objective literals p , $\neg p$ and their default negations. However, the added rules lead to undesired side effects that stand in direct opposition with basic principles underlying updates. Specifically, despite the fact that the empty program does not encode any change in the modelled world, the stable models assigned to a program may change after an update by the empty program.

This undesired behaviour is addressed in an alternative transformation from (Leite 2003) that encodes the coherence principle more carefully. Nevertheless, this transformation also leads to undesired consequences, as demonstrated in the following example:

Example 1.2 (Faulty sensor). *Suppose that we collect data from sensors and, for security reasons, multiple sensors are used to supply information about the critical fluent p . In case of a malfunction of one of the sensors, we may end up with an inconsistent logic program consisting of the following two facts:*

$$p. \quad \neg p.$$

At this point, no stable model of the program exists and action needs to be taken to find out what is wrong. If a problem is found in the sensor that supplied the first fact (p), after the sensor is repaired, this information needs to be reset by updating the program with the fact ($\sim p$). Following the universal pattern in rule updates, where recovery from conflicting states is always possible, we expect that this update is sufficient to assign a stable model to the updated program. However, the transformational semantics for strong negation defined in (Leite 2003) still does not provide any stable model – we remain without a valid epistemic state when one should in fact exist.

In this paper we address the issues with combining strong and default negation in the context of rule updates. Based on the above considerations, we formulate a generic desirable principle that is violated by the existing approaches. Then we show how two distinct definitions of one of the most well-behaved rule update semantics (Alferes et al. 2005; Banti et al. 2005) can be equivalently extended with support for strong negation. The resulting semantics not only satisfies the formulated principle, but also retains the formal and computational properties of the original semantics. More specifically, our main contributions are as follows:

- based on Example 1.2, we introduce the *early recovery principle* that captures circumstances under which a stable model after a rule update should exist;
- we extend the *well-supported semantics for rule updates* (Banti et al. 2005) with direct support for strong negation;
- we define a fixpoint characterisation of the new semantics, based on the *refined dynamic stable model* semantics for rule updates (Alferes et al. 2005);

- we show that the defined semantics enjoy the early recovery principle as well as a range of desirable properties for rule updates known from the literature.

This paper is organised as follows: In Sect. 2 we present the syntax and semantics of logic programs, generalise the well-supported semantics from the class of normal programs to extended ones and define the rule update semantics from (Alferes et al. 2005; Banti et al. 2005). Then, in Sect. 3, we formally establish the early recovery principle, define the new rule update semantics for strong negation and show that it satisfies the principle. In Sect. 4 we introduce other established rule update principles and show that the proposed semantics satisfies them. We discuss our findings and conclude in Sect. 5.¹

2 Background

In this section we introduce the necessary technical background and generalise the well-supported semantics (Fages 1991) to the class of extended programs.

2.1 Logic Programs

In the following we present the syntax of non-disjunctive logic programs with both strong and default negation in heads and bodies of rules, along with the definition of stable models of such programs from (Leite 2003) that is equivalent to the original definitions based on reducts (Gelfond and Lifschitz 1988; 1991; Inoue and Sakama 1998). Furthermore, we define an alternative characterisation of the stable model semantics: the well-supported models of normal logic programs (Fages 1991).

We assume that a countable set of propositional atoms \mathcal{A} is given and fixed. An *objective literal* is an atom $p \in \mathcal{A}$ or its strong negation $\neg p$. We denote the set of all objective literals by \mathcal{L} . A *default literal* is an objective literal preceded by \sim denoting default negation. A *literal* is either an objective or a default literal. We denote the set of all literals by \mathcal{L}^* . As a convention, double negation is absorbed, so that $\neg\neg p$ denotes the atom p and $\sim\sim l$ denotes the objective literal l . Given a set of literals S , we introduce the following notation: $S^+ = \{l \in \mathcal{L} \mid l \in S\}$, $S^- = \{l \in \mathcal{L} \mid \sim l \in S\}$, $\sim S = \{\sim l \mid l \in S\}$.

An *extended rule* is a pair $\pi = (H_\pi, B_\pi)$ where H_π is a literal, referred to as the *head* of π , and B_π is a finite set of literals, referred to as the *body* of π . Usually we write π as $(H_\pi \leftarrow B_\pi^+, \sim B_\pi^-)$. A *generalised rule* is an extended rule that contains no occurrence of \neg , i.e., its head and body consist only of atoms and their default negations. A *normal rule* is a generalised rule that has an atom in the head. A *fact* is an extended rule whose body is empty and a *tautology* is any extended rule π such that $H_\pi \in B_\pi$. An *extended (generalised, normal) program* is a set of extended (generalised, normal) rules.

An *interpretation* is a consistent subset of the set of objective literals, i.e., a subset of \mathcal{L} does not contain both p and $\neg p$ for any atom p . The satisfaction of an objective literal l ,

¹An extended version of this paper with all the proofs is available as (Slota, Baláž, and Leite 2014).

default literal $\sim l$, set of literals S , extended rule π and extended program P in an interpretation J is defined in the usual way: $J \models l$ iff $l \in J$; $J \models \sim l$ iff $l \notin J$; $J \models S$ iff $J \models L$ for all $L \in S$; $J \models \pi$ iff $J \models \mathbf{B}_\pi$ implies $J \models \mathbf{H}_\pi$; $J \models P$ iff $J \models \pi$ for all $\pi \in P$. Also, J is a *model* of P if $J \models P$, and P is *consistent* if it has a model.

Definition 2.1 (Stable model). *Let P be an extended program. The set $\llbracket P \rrbracket_{\text{SM}}$ of stable models of P consists of all interpretations J such that*

$$J^* = \text{least}(P \cup \text{def}(J))$$

where $\text{def}(J) = \{ \sim l. \mid l \in \mathcal{L} \setminus J \}$, $J^* = J \cup \sim(\mathcal{L} \setminus J)$ and $\text{least}(\cdot)$ denotes the least model of the argument program in which all literals are treated as propositional atoms.

A *level mapping* is a function that maps every atom to a natural number. Also, for any default literal $\sim p$, where $p \in \mathcal{A}$, and finite set of atoms and their default negations S , $\ell(\sim p) = \ell(p)$, $\ell^\downarrow(S) = \min \{ \ell(L) \mid L \in S \}$ and $\ell^\uparrow(S) = \max \{ \ell(L) \mid L \in S \}$.

Definition 2.2 (Well-supported model of a normal program). *Let P be a normal program and ℓ a level mapping. An interpretation $J \subseteq \mathcal{A}$ is a well-supported model of P w.r.t. ℓ if the following conditions are satisfied:*

1. J is a model of P ;
2. For every atom $p \in J$ there exists a rule $\pi \in P$ such that

$$\mathbf{H}_\pi = p \wedge J \models \mathbf{B}_\pi \wedge \ell(\mathbf{H}_\pi) > \ell^\uparrow(\mathbf{B}_\pi) .$$

The set $\llbracket P \rrbracket_{\text{WS}}$ of well-supported models of P consists of all interpretations $J \subseteq \mathcal{A}$ such that J is a well-supported model of P w.r.t. some level mapping.

As shown in (Fages 1991), well-supported models coincide with stable models:

Proposition 2.3 ((Fages 1991)). *Let P be a normal program. Then, $\llbracket P \rrbracket_{\text{WS}} = \llbracket P \rrbracket_{\text{SM}}$.*

2.2 Well-supported Models for Extended Programs

The well-supported models defined in the previous section for normal logic programs can be generalised in a straightforward manner to deal with strong negation while maintaining their tight relationship with stable models (c.f. Proposition 2.3). This will come useful in Subsect. 2.3 and Sect. 3 when we discuss adding support for strong negation to semantics for rule updates.

We extend level mappings from atoms and their default negations to all literals: An (*extended*) *level mapping* ℓ maps every objective literal to a natural number. Also, for any default literal $\sim l$ and finite set of literals S , $\ell(\sim l) = \ell(p)$, $\ell^\downarrow(S) = \min \{ \ell(L) \mid L \in S \}$ and $\ell^\uparrow(S) = \max \{ \ell(L) \mid L \in S \}$.

Definition 2.4 (Well-supported model of an extended program). *Let P be an extended program and ℓ a level mapping. An interpretation J is a well-supported model of P w.r.t. ℓ if the following conditions are satisfied:*

1. J is a model of P ;

2. For every objective literal $l \in J$ there exists a rule $\pi \in P$ such that

$$\mathbf{H}_\pi = l \wedge J \models \mathbf{B}_\pi \wedge \ell(\mathbf{H}_\pi) > \ell^\uparrow(\mathbf{B}_\pi) .$$

The set $\llbracket P \rrbracket_{\text{WS}}$ of well-supported models of P consists of all interpretations J such that J is a well-supported model of P w.r.t. some level mapping.

We obtain a generalisation of Prop. 2.3 to the class of extended programs:

Proposition 2.5. *Let P be an extended program. Then, $\llbracket P \rrbracket_{\text{WS}} = \llbracket P \rrbracket_{\text{SM}}$.*

2.3 Rule Updates

We turn our attention to rule updates, starting with one of the most advanced rule update semantics, the *refined dynamic stable models* for sequences of generalised programs (Alferes et al. 2005), as well as the equivalent definition of *well-supported models* (Banti et al. 2005). Then we define the transformations for adding support for strong negation to such semantics (Alferes and Pereira 1996; Leite 2003).

A rule update semantics provides a way to assign stable models to a pair or sequence of programs where each component represents an update of the preceding ones. Formally, a *dynamic logic program* (DLP) is a finite sequence of extended programs and by $\text{all}(\mathbf{P})$ we denote the multiset of all rules in the components of \mathbf{P} . A rule update semantics \mathbf{S} assigns a *set of S-models*, denoted by $\llbracket \mathbf{P} \rrbracket_{\mathbf{S}}$, to \mathbf{P} .

We focus on semantics based on the causal rejection principle (Leite and Pereira 1998; Alferes et al. 2000; Eiter et al. 2002; Leite 2003; Alferes et al. 2005; Banti et al. 2005; Osorio and Cuevas 2007) which states that a rule is *rejected* if it is in a direct conflict with a more recent rule. The basic type of conflict between rules π and σ occurs when their heads contain complementary literals, i.e. when $\mathbf{H}_\pi = \sim \mathbf{H}_\sigma$. Based on such conflicts and on a stable model candidate, a *set of rejected rules* can be determined and it can be verified that the candidate is indeed stable w.r.t. the remaining rules.

We define the most mature of these semantics, providing two equivalent definitions: the *refined dynamic stable models* (Alferes et al. 2005), or *RD-semantics*, defined using a fixpoint equation, and the *well-supported models* (Banti et al. 2005), or *WS-semantics*, based on level mappings.

Definition 2.6 (RD-semantics (Alferes et al. 2005)). *Let $\mathbf{P} = \langle P_i \rangle_{i < n}$ be a DLP without strong negation. Given an interpretation J , the multisets of rejected rules $\text{rej}_{\geq}(\mathbf{P}, J)$ and of default assumptions $\text{def}(\mathbf{P}, J)$ are defined as follows:*

$$\text{rej}_{\geq}(\mathbf{P}, J) = \{ \pi \in P_i \mid i < n \wedge \exists j \geq i \exists \sigma \in P_j : \mathbf{H}_\pi = \sim \mathbf{H}_\sigma \wedge J \models \mathbf{B}_\sigma \},$$

$$\text{def}(\mathbf{P}, J) = \{ (\sim l.) \mid l \in \mathcal{L} \wedge \neg(\exists \pi \in \text{all}(\mathbf{P}) : \mathbf{H}_\pi = l \wedge J \models \mathbf{B}_\pi) \}.$$

The set $\llbracket \mathbf{P} \rrbracket_{\text{RD}}$ of RD-models of \mathbf{P} consists of all interpretations J such that

$$J^* = \text{least}([\text{all}(\mathbf{P}) \setminus \text{rej}_{\geq}(\mathbf{P}, J)] \cup \text{def}(\mathbf{P}, J))$$

where J^* and $\text{least}(\cdot)$ are defined as before.

Definition 2.7 (WS-semantics (Banti et al. 2005)). Let $\mathbf{P} = \langle P_i \rangle_{i < n}$ be a DLP without strong negation. Given an interpretation J and a level mapping ℓ , the multiset of rejected rules $\text{rej}_\ell(\mathbf{P}, J)$ is defined as follows:

$$\text{rej}_\ell(\mathbf{P}, J) = \{ \pi \in P_i \mid i < n \ \exists j > i \ \exists \sigma \in P_j : H_\pi = \sim H_\sigma \wedge J \models B_\sigma \wedge \ell(H_\sigma) > \ell^\dagger(B_\sigma) \}.$$

The set $\llbracket \mathbf{P} \rrbracket_{\text{WS}}$ of WS-models of \mathbf{P} consists of all interpretations J such that for some level mapping ℓ , the following conditions are satisfied:

1. J is a model of $\text{all}(\mathbf{P}) \setminus \text{rej}_\ell(\mathbf{P}, J)$;
2. For every $l \in J$ there exists some rule $\pi \in \text{all}(\mathbf{P}) \setminus \text{rej}_\ell(\mathbf{P}, J)$ such that

$$H_\pi = l \wedge J \models B_\pi \wedge \ell(H_\pi) > \ell^\dagger(B_\pi).$$

Unlike most other rule update semantics, these semantics can properly deal with tautological and other irrelevant updates, as illustrated in the following example:

Example 2.8 (Irrelevant updates). Consider the DLP $\mathbf{P} = \langle P, U \rangle$ where programs P, U are as follows:

$$\begin{aligned} P : \quad & \text{day} \leftarrow \sim \text{night}. & \text{stars} \leftarrow \text{night}, \sim \text{cloudy}. \\ & \text{night} \leftarrow \sim \text{day}. & \sim \text{stars}. \\ U : \quad & \text{stars} \leftarrow \text{stars}. \end{aligned}$$

Note that program P has the single stable model $J_1 = \{ \text{day} \}$ and U contains a single tautological rule, i.e. it does not encode any change in the modelled domain. Thus, we expect that \mathbf{P} also has the single stable model J_1 . Nevertheless, many rule update semantics, such as those introduced in (Leite and Pereira 1998; Alferes et al. 2000; Eiter et al. 2002; Leite 2003; Sakama and Inoue 2003; Zhang 2006; Osorio and Cuevas 2007; Delgrande, Schaub, and Tompits 2007; Krümpelmann 2012), are sensitive to this or other tautological updates, introducing or eliminating models of the original program.

In this case, the unwanted model candidate is $J_2 = \{ \text{night}, \text{stars} \}$ and it is neither an RD- nor a WS-model of \mathbf{P} , though the reasons for this are technically different under these two semantics. It is not difficult to verify that, given an arbitrary level mapping ℓ , the respective sets of rejected rules and the set of default assumptions are as follows:

$$\begin{aligned} \text{rej}_{\geq}(\mathbf{P}, J_2) &= \{ (\text{stars} \leftarrow \text{night}, \sim \text{cloudy}.), (\sim \text{stars}.), \\ \text{rej}_\ell(\mathbf{P}, J_2) &= \emptyset, \\ \text{def}(\mathbf{P}, J_2) &= \{ (\sim \text{cloudy}.), (\sim \text{day}.), \}. \end{aligned}$$

Note that $\text{rej}_\ell(\mathbf{P}, J_2)$ is empty because, independently of ℓ , no rule π in U satisfies the condition $\ell(H_\pi) > \ell^\dagger(B_\pi)$, so there is no rule that could reject another rule. Thus, the atom stars belongs to J_2^* but does not belong to $\text{least}(\llbracket \text{all}(\mathbf{P}) \setminus \text{rej}_{\geq}(\mathbf{P}, J_2) \rrbracket \cup \text{def}(\mathbf{P}, J_2))$, so J_2 is not an RD-model of \mathbf{P} . Furthermore, no model of $\text{all}(\mathbf{P}) \setminus \text{rej}_\ell(\mathbf{P}, J_2)$ contains stars , so J_2 cannot be a WS-model of \mathbf{P} .

Furthermore, the resilience of RD- and WS-semantics is not limited to empty and tautological updates, but extends to other irrelevant updates as well (Alferes et al. 2005;

Banti et al. 2005). For example, consider the DLP $\mathbf{P}' = \langle P, U' \rangle$ where $U' = \{ (\text{stars} \leftarrow \text{venus}.), (\text{venus} \leftarrow \text{stars}.), \}$. Though the updating program contains non-tautological rules, it does not provide a bottom-up justification of any model other than J_1 and, indeed, J_1 is the only RD- and WS-model of \mathbf{P}' .

We also note that the two presented semantics for DLPs without strong negation provide the same result regardless of the particular DLP to which they are applied.

Proposition 2.9 ((Banti et al. 2005)). Let \mathbf{P} be a DLP without strong negation. Then, $\llbracket \mathbf{P} \rrbracket_{\text{WS}} = \llbracket \mathbf{P} \rrbracket_{\text{RD}}$.

In case of the stable model semantics for a single program, strong negation can be reduced away by treating all objective literals as atoms and adding, for each atom p , the integrity constraint $(\leftarrow p, \neg p)$ to the program (Gelfond and Lifschitz 1991). However, this transformation does not serve its purpose when adding support for strong negation to causal rejection semantics for DLPs because integrity constraints have empty heads, so according to these rule update semantics, they cannot be used to reject any other rule. For example, a DLP such as $\langle \{ p., \neg p. \}, \{ p. \} \rangle$ would remain without a stable model even though the DLP $\langle \{ p., \sim p. \}, \{ p. \} \rangle$ does have a stable model.

To capture the conflict between opposite objective literals l and $\neg l$ in a way that is compatible with causal rejection semantics, a slightly modified syntactic transformation can be performed, translating such conflicts into conflicts between objective literals and their default negations. Two such transformations have been suggested in the literature (Alferes and Pereira 1996; Leite 2003), both based on the principle of coherence. For any extended program P and DLP $\mathbf{P} = \langle P_i \rangle_{i < n}$ they are defined as follows:

$$\begin{aligned} P^\dagger &= P \cup \{ \sim \neg l \leftarrow l. \mid l \in \mathcal{L} \}, \\ \mathbf{P}^\dagger &= \left\langle P_i^\dagger \right\rangle_{i < n}, \\ P^\ddagger &= P \cup \{ \sim \neg H_\pi \leftarrow B_\pi. \mid \pi \in P \wedge H_\pi \in \mathcal{L} \}, \\ \mathbf{P}^\ddagger &= \left\langle P_i^\ddagger \right\rangle_{i < n}. \end{aligned}$$

These transformations lead to four possibilities for defining the semantics of an arbitrary DLP \mathbf{P} : $\llbracket \mathbf{P}^\dagger \rrbracket_{\text{RD}}$, $\llbracket \mathbf{P}^\ddagger \rrbracket_{\text{RD}}$, $\llbracket \mathbf{P}^\dagger \rrbracket_{\text{WS}}$ and $\llbracket \mathbf{P}^\ddagger \rrbracket_{\text{WS}}$. We discuss these in the following section.

3 Direct Support for Strong Negation in Rule Updates

The problem with existing semantics for strong negation in rule updates is that semantics based on the first transformation (\mathbf{P}^\dagger) assign too many models to some DLPs, while semantics based on the second transformation (\mathbf{P}^\ddagger) sometimes do not assign any model to a DLP that should have one. The former is illustrated in the following example:

Example 3.1 (Undesired side effects of the first transformation). Consider the DLP $\mathbf{P}_1 = \langle P, U \rangle$ where $P = \{ p., \neg p. \}$ and $U = \emptyset$. Since P has no stable model and U does not encode any change in the represented domain, it should follow that \mathbf{P}_1 has no stable model either. However, $\llbracket \mathbf{P}_1^\dagger \rrbracket_{\text{RD}} =$

$\llbracket \mathbf{P}_1^\dagger \rrbracket_{\text{WS}} = \{ \{p\}, \{\neg p\} \}$, i.e. two models are assigned to \mathbf{P}_1 when using the first transformation to add support for strong negation. To verify this, observe that $\mathbf{P}_1^\dagger = \langle P^\dagger, U^\dagger \rangle$ where

$$\begin{array}{llll} P^\dagger : & p. & \neg p. & U^\dagger : \quad \sim p \leftarrow \neg p. \\ & \sim p \leftarrow \neg p. & \sim \neg p \leftarrow p. & \sim \neg p \leftarrow p. \end{array}$$

Consider the interpretation $J_1 = \{p\}$. It is not difficult to verify that

$$\begin{aligned} \text{rej}_{\geq}(\mathbf{P}_1^\dagger, J_1) &= \{ \neg p., \sim \neg p \leftarrow p. \} , \\ \text{def}(\mathbf{P}_1^\dagger, J_1) &= \emptyset , \end{aligned}$$

so it follows that

$$\begin{aligned} \text{least} \left(\left[\text{all}(\mathbf{P}_1^\dagger) \setminus \text{rej}_{\geq}(\mathbf{P}_1^\dagger, J_1) \right] \cup \text{def}(\mathbf{P}_1^\dagger, J_1) \right) &= \\ &= \{p, \sim \neg p\} = J_1^* . \end{aligned}$$

In other words, J_1 belongs to $\llbracket \mathbf{P}_1^\dagger \rrbracket_{\text{RD}}$ and in an analogous fashion it can be verified that $J_2 = \{\neg p\}$ also belongs there. A similar situation occurs with $\llbracket \mathbf{P}_1^\dagger \rrbracket_{\text{WS}}$ since the rules that were added to the more recent program can be used to reject facts in the older one.

Thus, the problem with the first transformation is that an update by an empty program, which does not express any change in the represented domain, may affect the original semantics. This behaviour goes against basic and intuitive principles underlying updates, grounded already in the classical belief update postulates (Keller and Winslett 1985; Katsuno and Mendelzon 1991) and satisfied by virtually all belief update operations (Herzig and Rifi 1999) as well as by the vast majority of existing rule update semantics, including the original RD- and WS-semantics.

This undesired behaviour can be corrected by using the second transformation instead. The more technical reason is that it does not add any rules to a program in the sequence unless that program already contains some original rules. However, its use leads to another problem: sometimes *no model* is assigned when in fact a model should exist.

Example 3.2 (Undesired side effects of the second transformation). Consider again Example 1.2, formalised as the DLP $\mathbf{P}_2 = \langle P, V \rangle$ where $P = \{p., \neg p.\}$ and $V = \{\sim p.\}$. It is reasonable to expect that since V resolves the conflict present in P , a stable model should be assigned to \mathbf{P}_2 . However, $\llbracket \mathbf{P}_2^\dagger \rrbracket_{\text{RD}} = \llbracket \mathbf{P}_2^\dagger \rrbracket_{\text{WS}} = \emptyset$. To verify this, observe that $\mathbf{P}_2^\dagger = \langle P^\dagger, V^\dagger \rangle$ where

$$\begin{array}{llll} P^\dagger : & p. & \neg p. & V^\dagger : \quad \sim p. \\ & \sim p. & \sim \neg p. & \end{array}$$

Given an interpretation J and level mapping ℓ , we conclude that $\text{rej}_\ell(\mathbf{P}_2^\dagger, J) = \{p.\}$, so the facts $(\neg p.)$ and $(\sim \neg p.)$ both belong to the program

$$\text{all}(\mathbf{P}_2^\dagger) \setminus \text{rej}_\ell(\mathbf{P}_2^\dagger, J) .$$

Consequently, this program has no model and it follows that J cannot belong to $\llbracket \mathbf{P}_2^\dagger \rrbracket_{\text{WS}}$. Similarly it can be shown that $\llbracket \mathbf{P}_2^\dagger \rrbracket_{\text{RD}} = \emptyset$.

Based on this example, in the following we formulate a generic *early recovery principle* that formally identifies conditions under which *some* stable model should be assigned to a DLP. For the sake of simplicity, we concentrate on DLPs of length 2 which are composed of facts. We discuss a generalisation of the principle to DLPs of arbitrary length and containing other rules than just facts in Sect. 5. After introducing the principle, we define a semantics for rule updates which directly supports both strong and default negation and satisfies the principle.

We begin by defining, for every objective literal l , the sets of literals \bar{l} and $\overline{\sim}l$ as follows:

$$\bar{l} = \{ \sim l, \neg l \} \quad \text{and} \quad \overline{\sim}l = \{ l \} .$$

Intuitively, for every literal L , \bar{L} denotes the set of literals that are in conflict with L . Furthermore, given two sets of facts P and U , we say that U *solves all conflicts in* P if for each pair of rules $\pi, \sigma \in P$ such that $H_\sigma \in \overline{H_\pi}$ there is a fact $\rho \in U$ such that either $H_\rho \in \overline{H_\pi}$ or $H_\rho \in \overline{H_\sigma}$.

Considering a rule update semantics S , the new principle simply requires that when U solves all conflicts in P , S will assign *some model* to $\langle P, U \rangle$. Formally:

Early recovery principle: If P is a set of facts and U is a consistent set of facts that solves all conflicts in P , then $\llbracket \langle P, U \rangle \rrbracket_S \neq \emptyset$.

We conjecture that rule update semantics should generally satisfy the above principle. In contrast with the usual behaviour of belief update operators, the nature of existing rule update semantics ensures that recovery from conflict is always possible, and this principle simply formalises and sharpens the sufficient conditions for such recovery.

Our next goal is to define a semantics for rule updates that not only satisfies the outlined principle, but also enjoys other established properties of rule updates that have been identified over the years. Similarly as for the original semantics for rule updates, we provide two equivalent definitions, one based on a fixed point equation and the other one on level mappings.

To directly accommodate strong negation in the RD-semantics, we first need to look more closely at the set of rejected rules $\text{rej}_{\geq}(\mathbf{P}, J)$, particularly at the fact that it allows conflicting rules within the same component of \mathbf{P} to reject one another. This behaviour, along with the constrained set of defaults $\text{def}(\mathbf{P}, J)$, is used to prevent tautological and other irrelevant cyclic updates from affecting the semantics. However, in the presence of strong negation, rejecting conflicting rules within the same program has undesired side effects. For example, the early recovery principle requires that some model be assigned to the DLP $\langle \{p., \neg p.\}, \{\sim p.\} \rangle$ from Example 3.2, but if the rules in the initial program reject each other, then the only possible stable model to assign is \emptyset . However, such a stable model would violate the causal rejection principle since it does not satisfy the initial rule $(\neg p.)$ and there is no rule in the updating program that overrides it.

To overcome the limitations of this approach to the prevention of tautological updates, we disentangle rule rejection per se from ensuring that rejection is done without

cyclic justifications. We introduce the set of rejected rules $\text{rej}_{>}^{\neg}(\mathbf{P}, S)$ which directly supports strong negation and does not allow for rejection within the same program. Prevention of cyclic rejections is done separately by using a customised immediate consequence operator $T_{\mathbf{P}, J}$. Given a stable model candidate J , instead of verifying that J^* is the least fixed point of the usual consequence operator, as done in the RD-semantics using $\text{least}(\cdot)$, we verify that J^* is the least fixed point of $T_{\mathbf{P}, J}$.

Definition 3.3 (Extended RD-semantics). *Let $\mathbf{P} = \langle P_i \rangle_{i < n}$ be a DLP. Given an interpretation J and a set of literals S , the multiset of rejected rules $\text{rej}_{>}^{\neg}(\mathbf{P}, S)$, the remainder $\text{rem}(\mathbf{P}, S)$ and the consequence operator $T_{\mathbf{P}, J}$ are defined as follows:*

$$\text{rej}_{>}^{\neg}(\mathbf{P}, S) = \{ \pi \in P_i \mid i < n \wedge \exists j > i \exists \sigma \in P_j : \text{H}_{\sigma} \in \overline{\text{H}_{\pi}} \wedge \text{B}_{\sigma} \subseteq S \},$$

$$\text{rem}(\mathbf{P}, S) = \text{all}(\mathbf{P}) \setminus \text{rej}_{>}^{\neg}(\mathbf{P}, S) ,$$

$$T_{\mathbf{P}, J}(S) = \{ \text{H}_{\pi} \mid \pi \in (\text{rem}(\mathbf{P}, J^*) \cup \text{def}(J)) \wedge \text{B}_{\pi} \subseteq S \wedge \neg (\exists \sigma \in \text{rem}(\mathbf{P}, S) : \text{H}_{\sigma} \in \overline{\text{H}_{\pi}} \wedge \text{B}_{\sigma} \subseteq J^*) \} .$$

Furthermore, $T_{\mathbf{P}, J}^0(S) = S$ and for every $k \geq 0$, $T_{\mathbf{P}, J}^{k+1}(S) = T_{\mathbf{P}, J}(T_{\mathbf{P}, J}^k(S))$. The set $\llbracket \mathbf{P} \rrbracket_{\text{RD}}^{\neg}$ of extended RD-models of \mathbf{P} consists of all interpretations J such that

$$J^* = \bigcup_{k \geq 0} T_{\mathbf{P}, J}^k(\emptyset) .$$

Adding support for strong negation to the WS-semantics is done by modifying the set of rejected rules $\text{rej}_{\ell}(\mathbf{P}, J)$ to account for the new type of conflict. Additionally, in order to ensure that rejection of a literal L cannot be based on the assumption that some conflicting literal $L' \in \overline{L}$ is true, a rejecting rule σ must satisfy the stronger condition $\ell^{\downarrow}(\overline{L}) > \ell^{\uparrow}(\text{B}_{\sigma})$. Finally, to prevent defeated rules from affecting the resulting models, we require that all supporting rules belong to $\text{rem}(\mathbf{P}, J^*)$.

Definition 3.4 (Extended WS-semantics). *Let $\mathbf{P} = \langle P_i \rangle_{i < n}$ be a DLP. Given an interpretation J and a level mapping ℓ , the multiset of rejected rules $\text{rej}_{\ell}^{\neg}(\mathbf{P}, J)$ is defined by:*

$$\text{rej}_{\ell}^{\neg}(\mathbf{P}, J) = \{ \pi \in P_i \mid i < n \wedge \exists j > i \exists \sigma \in P_j : \text{H}_{\sigma} \in \overline{\text{H}_{\pi}} \wedge J \models \text{B}_{\sigma} \wedge \ell^{\downarrow}(\overline{\text{H}_{\pi}}) > \ell^{\uparrow}(\text{B}_{\sigma}) \} .$$

The set $\llbracket \mathbf{P} \rrbracket_{\text{WS}}^{\neg}$ of extended WS-models of \mathbf{P} consists of all interpretations J such that for some level mapping ℓ , the following conditions are satisfied:

1. J is a model of $\text{all}(\mathbf{P}) \setminus \text{rej}_{\ell}^{\neg}(\mathbf{P}, J)$;
2. For every $l \in J$ there exists some rule $\pi \in \text{rem}(\mathbf{P}, J^*)$ such that

$$\text{H}_{\pi} = l \wedge J \models \text{B}_{\pi} \wedge \ell(\text{H}_{\pi}) > \ell^{\uparrow}(\text{B}_{\pi}) .$$

The following theorem establishes that the two defined semantics are equivalent:

Theorem 3.5. *Let \mathbf{P} be a DLP. Then, $\llbracket \mathbf{P} \rrbracket_{\text{WS}}^{\neg} = \llbracket \mathbf{P} \rrbracket_{\text{RD}}^{\neg}$.*

Also, on DLPs without strong negation they coincide with the original semantics.

Theorem 3.6. *Let \mathbf{P} be a DLP without strong negation. Then, $\llbracket \mathbf{P} \rrbracket_{\text{WS}}^{\neg} = \llbracket \mathbf{P} \rrbracket_{\text{RD}}^{\neg} = \llbracket \mathbf{P} \rrbracket_{\text{WS}} = \llbracket \mathbf{P} \rrbracket_{\text{RD}}$.*

Furthermore, unlike the transformational semantics for strong negation, the new semantics satisfy the early recovery principle.

Theorem 3.7. *The extended RD-semantics and extended WS-semantics satisfy the early recovery principle.*

4 Properties

In this section we take a closer look at the formal and computational properties of the proposed rule update semantics.

The various approaches to rule updates (Leite and Pereira 1998; Alferes et al. 2000; Eiter et al. 2002; Leite 2003; Sakama and Inoue 2003; Alferes et al. 2005; Banti et al. 2005; Zhang 2006; Šefranek 2006; Osorio and Cuevas 2007; Delgrande, Schaub, and Tompits 2007; Šefranek 2011; Krumpelmann 2012) share a number of basic characteristics. For example, all of them generalise stable models, i.e., the models they assign to a sequence $\langle P \rangle$ (of length 1) are exactly the stable models of P . Similarly, they adhere to the principle of primacy of new information (Dalal 1988), so models assigned to $\langle P_i \rangle_{i < n}$ satisfy the latest program P_{n-1} . However, they also differ significantly in their technical realisation and classes of supported inputs, and desirable properties such as immunity to tautologies are violated by many of them.

Table 1 lists many of the generic properties proposed for rule updates that have been identified and formalised throughout the years (Leite and Pereira 1998; Eiter et al. 2002; Leite 2003; Alferes et al. 2005). The rule update semantics we defined in the previous section enjoys all of them.

Theorem 4.1. *The extended RD-semantics and extended WS-semantics satisfy all properties listed in Table 1.*

Our semantics also retains the same computational complexity as the stable models.

Theorem 4.2. *Let \mathbf{P} be a DLP. The problem of deciding whether some $J \in \llbracket \mathbf{P} \rrbracket_{\text{WS}}^{\neg}$ exists is NP-complete. Given a literal L , the problem of deciding whether for all $J \in \llbracket \mathbf{P} \rrbracket_{\text{WS}}^{\neg}$ it holds that $J \models L$ is coNP-complete.*

5 Concluding Remarks

In this paper we have identified shortcomings in the existing semantics for rule updates that fully support both strong and default negation, and proposed a generic *early recovery principle* that captures them formally. Subsequently, we provided two equivalent definitions of a new semantics for rule updates.

We have shown that the newly introduced rule update semantics constitutes a strict improvement upon the state of the art in rule updates as it enjoys the following combination of characteristics, unmatched by any previously existing semantics:

- It allows for both strong and default negation in heads of rules, making it possible to move between any pair of epistemic states by means of updates;

Table 1: Desirable properties of rule update semantics

Generalisation of stable models	$\llbracket \langle P \rangle \rrbracket_s = \llbracket P \rrbracket_{SM}$.
Primacy of new information	If $J \in \llbracket \langle P_i \rangle_{i < n} \rrbracket_s$, then $J \models P_{n-1}$.
Fact update	A sequence of consistent sets of facts $\langle P_i \rangle_{i < n}$ has the single model $\{ l \in \mathcal{L} \mid \exists i < n : (l.) \in P_i \wedge (\forall j > i : \{ \neg l., \sim l. \} \cap P_j = \emptyset) \}$.
Support	If $J \in \llbracket P \rrbracket_s$ and $l \in J$, then there is some rule $\pi \in \text{all}(P)$ such that $H_\pi = l$ and $J \models B_\pi$.
Idempotence	$\llbracket \langle P, P \rangle \rrbracket_s = \llbracket \langle P \rangle \rrbracket_s$.
Absorption	$\llbracket \langle P, U, U \rangle \rrbracket_s = \llbracket \langle P, U \rangle \rrbracket_s$.
Augmentation	If $U \subseteq V$, then $\llbracket \langle P, U, V \rangle \rrbracket_s = \llbracket \langle P, V \rangle \rrbracket_s$.
Non-interference	If U and V are over disjoint alphabets, then $\llbracket \langle P, U, V \rangle \rrbracket_s = \llbracket \langle P, V, U \rangle \rrbracket_s$.
Immunity to empty updates	If $P_j = \emptyset$, then $\llbracket \langle P_i \rangle_{i < n} \rrbracket_s = \llbracket \langle P_i \rangle_{i < n \wedge i \neq j} \rrbracket_s$.
Immunity to tautologies	If $\langle Q_i \rangle_{i < n}$ is a sequence of sets of tautologies, then $\llbracket \langle P_i \cup Q_i \rangle_{i < n} \rrbracket_s = \llbracket \langle P_i \rangle_{i < n} \rrbracket_s$.
Causal rejection principle	For every $i < n$, $\pi \in P_i$ and $J \in \llbracket \langle P_i \rangle_{i < n} \rrbracket_s$, if $J \not\models \pi$, then there exists some $\sigma \in P_j$ with $j > i$ such that $H_\sigma \in \overline{H_\pi}$ and $J \models B_\sigma$.

- It satisfies the *early recovery principle* which guarantees the existence of a model whenever all conflicts in the original program are satisfied;
- It enjoys all rule update principles and desirable properties reported in Table 1;
- It does not increase the computational complexity of the stable model semantics upon which it is based.

However, the early recovery principle, as it is formulated in Sect. 3, only covers a single update of a set of facts by another set of facts. Can it be generalised further without rendering it too strong? Certain caution is appropriate here, since in general the absence of a stable model can be caused by odd cycles or simply by the fundamental differences between different approaches to rule update, and the purpose of this principle is not to choose which approach to take.

Nevertheless, one generalisation that should cause no harm is the generalisation to iterated updates, i.e. to sequences of sets of facts. Another generalisation that appears very reasonable is the generalisation to *acyclic DLPs*, i.e. DLPs such that $\text{all}(P)$ is an acyclic program. An acyclic program has at most one stable model, and if we guarantee that all potential conflicts within it certainly get resolved, we can safely conclude that the rule update semantics should assign some model to it. We formalise these ideas in what follows.

We say that a program P is *acyclic* (Apt and Bezem 1991) if for some level mapping ℓ , such that for every $l \in \mathcal{L}$, $\ell(l) = \ell(\neg l)$, and every rule $\pi \in P$ it holds that $\ell(H_\pi) > \ell^\uparrow(B_\pi)$. Given a DLP $P = \langle P_i \rangle_{i < n}$, we say that *all conflicts in P are solved* if for every $i < n$ and each pair of rules $\pi, \sigma \in P_i$ such that $H_\sigma \in \overline{H_\pi}$ there is some $j > i$ and a fact $\rho \in P_j$ such that either $H_\rho \in \overline{H_\pi}$ or $H_\rho \in \overline{H_\sigma}$.

Generalised early recovery principle: If $\text{all}(P)$ is acyclic and all conflicts in P are solved, then $\llbracket P \rrbracket_s \neq \emptyset$.

Note that this generalisation of the early recovery princi-

ple applies to a much broader class of DLPs than the original one. We illustrate this in the following example:

Example 5.1 (Recovery in a stratified program). *Consider the following programs P, U and V :*

$$\begin{array}{l}
P : \quad p \leftarrow q, \sim r. \quad \sim p \leftarrow s. \quad q. \quad s \leftarrow q. \\
U : \quad \neg p. \quad r \leftarrow q. \quad \neg r \leftarrow q, s. \\
V : \quad \sim r.
\end{array}$$

Looking more closely at program P , we see that atoms q and s are derived by the latter two rules inside it while atom r is false by default since there is no rule that could be used to derive its truth. Consequently, the bodies of the first two rules are both satisfied and as their heads are conflicting, P has no stable model. The single conflict in P is solved after it is updated by U , but then another conflict is introduced due to the latter two rules in the updating program. This second conflict can be solved after another update by V . Consequently, we expect that some stable model be assigned to the DLP $\langle P, U, V \rangle$.

The original early recovery principle does not impose this because the DLP in question has more than two components and the rules within it are not only facts. However, the DLP is acyclic, as shown by any level mapping ℓ with $\ell(p) = 3$, $\ell(q) = 0$, $\ell(r) = 2$ and $\ell(s) = 1$, so the generalised early recovery principle does apply. Furthermore, we also find the single extended RD-model of $\langle P, U, V \rangle$ is $\{ \neg p, q, \neg r, s \}$, i.e. the semantics respects the stronger principle in this case.

Moreover, as established in the following theorem, it is no coincidence that the extended RD-semantics respects the stronger principle in the above example – the principle is generally satisfied by the semantics introduced in this paper.

Theorem 5.2. *The extended RD-semantics and extended WS-semantics satisfy the generalised early recovery principle.*

Both the original and the generalised early recovery principle can guide the future addition of full support for both kinds of negations in other approaches to rule updates, such as those proposed in (Sakama and Inoue 2003; Zhang 2006; Delgrande, Schaub, and Tompits 2007; Krümpelmann 2012), making it possible to reach any belief state by updating the current program. Furthermore, adding support for strong negation is also interesting in the context of recent results on program revision and updates that are performed on the *semantic level*, ensuring syntax-independence of the respective methods (Delgrande et al. 2013; Slota and Leite 2014; 2012a; 2010), in the context of finding suitable condensing operators (Slota and Leite 2013), and unifying with updates in classical logic (Slota and Leite 2012b).

Acknowledgments

João Leite was partially supported by Fundação para a Ciência e a Tecnologia under project “ERRO – Efficient Reasoning with Rules and Ontologies” (PTDC/EIA-CCO/121823/2010). Martin Slota was partially supported by Fundação para a Ciência e a Tecnologia under project “ASPEN – Answer Set Programming with BoolEaN Satisfiability” (PTDC/EIA-CCO/110921/2009). The collaboration between the co-authors resulted from the Slovak–Portuguese bilateral project “ReDIK – Reasoning with Dynamic Inconsistent Knowledge”, supported by APVV agency under SK-PT-0028-10 and by Fundação para a Ciência e a Tecnologia (FCT/2487/3/6/2011/S).

References

- Alferes, J. J., and Pereira, L. M. 1996. Update-programs can update programs. In Dix, J.; Pereira, L. M.; and Przymusinski, T. C., eds., *Non-Monotonic Extensions of Logic Programming (NMELP '96), Selected Papers*, volume 1216 of *Lecture Notes in Computer Science*, 110–131. Bad Honnef, Germany: Springer.
- Alferes, J. J.; Leite, J. A.; Pereira, L. M.; Przymusinska, H.; and Przymusinski, T. C. 1998. Dynamic logic programming. In Cohn, A. G.; Schubert, L. K.; and Shapiro, S. C., eds., *Proceedings of the Sixth International Conference on Principles of Knowledge Representation and Reasoning (KR'98), Trento, Italy, June 2-5, 1998*, 98–111. Morgan Kaufmann.
- Alferes, J. J.; Leite, J. A.; Pereira, L. M.; Przymusinska, H.; and Przymusinski, T. C. 2000. Dynamic updates of non-monotonic knowledge bases. *The Journal of Logic Programming* 45(1-3):43–70.
- Alferes, J. J.; Banti, F.; Brogi, A.; and Leite, J. A. 2005. The refined extension principle for semantics of dynamic logic programming. *Studia Logica* 79(1):7–32.
- Apt, K. R., and Bezem, M. 1991. Acyclic programs. *New Generation Computing* 9(3/4):335–364.
- Banti, F.; Alferes, J. J.; Brogi, A.; and Hitzler, P. 2005. The well supported semantics for multidimensional dynamic logic programs. In Baral, C.; Greco, G.; Leone, N.; and Terracina, G., eds., *Proceedings of the 8th International Conference on Logic Programming and Nonmonotonic Reasoning (LPNMR 2005)*, volume 3662 of *Lecture Notes in Computer Science*, 356–368. Diamante, Italy: Springer.
- Berners-Lee, T.; Hendler, J.; and Lassila, O. 2001. The semantic web. *Scientific American* 284(5):28–37.
- Dalal, M. 1988. Investigations into a theory of knowledge base revision. In *Proceedings of the 7th National Conference on Artificial Intelligence (AAAI 1988)*, 475–479. St. Paul, MN, USA: AAAI Press / The MIT Press.
- Delgrande, J.; Schaub, T.; Tompits, H.; and Woltran, S. 2013. A model-theoretic approach to belief change in answer set programming. *ACM Transactions on Computational Logic (TOCL)* 14(2):14:1–14:46.
- Delgrande, J. P.; Schaub, T.; and Tompits, H. 2007. A preference-based framework for updating logic programs. In Baral, C.; Brewka, G.; and Schlipf, J. S., eds., *Proceedings of the 9th International Conference on Logic Programming and Nonmonotonic Reasoning (LPNMR 2007)*, volume 4483 of *Lecture Notes in Computer Science*, 71–83. Tempe, AZ, USA: Springer.
- Eiter, T.; Fink, M.; Sabbatini, G.; and Tompits, H. 2002. On properties of update sequences based on causal rejection. *Theory and Practice of Logic Programming (TPLP)* 2(6):721–777.
- Fages, F. 1991. A new fixpoint semantics for general logic programs compared with the well-founded and the stable model semantics. *New Generation Computing* 9(3/4):425–444.
- Gelfond, M., and Lifschitz, V. 1988. The stable model semantics for logic programming. In Kowalski, R. A., and Bowen, K. A., eds., *Proceedings of the 5th International Conference and Symposium on Logic Programming (ICLP/SLP 1988)*, 1070–1080. Seattle, Washington: MIT Press.
- Gelfond, M., and Lifschitz, V. 1991. Classical negation in logic programs and disjunctive databases. *New Generation Computing* 9(3-4):365–385.
- Herzig, A., and Rifi, O. 1999. Propositional belief base update and minimal change. *Artificial Intelligence* 115(1):107–138.
- Inoue, K., and Sakama, C. 1998. Negation as failure in the head. *Journal of Logic Programming* 35(1):39–78.
- Katsuno, H., and Mendelzon, A. O. 1991. On the difference between updating a knowledge base and revising it. In Allen, J. F.; Fikes, R.; and Sandewall, E., eds., *Proceedings of the 2nd International Conference on Principles of Knowledge Representation and Reasoning (KR'91)*, 387–394. Cambridge, MA, USA: Morgan Kaufmann Publishers.
- Keller, A. M., and Winslett, M. 1985. On the use of an extended relational model to handle changing incomplete information. *IEEE Transactions on Software Engineering* 11(7):620–633.
- Krümpelmann, P. 2012. Dependency semantics for sequences of extended logic programs. *Logic Journal of the IGPL* 20(5):943–966.
- Leite, J. A., and Pereira, L. M. 1998. Generalizing updates: From models to programs. In Dix, J.; Pereira, L. M.;

- and Przymusiński, T. C., eds., *Proceedings of the 3rd International Workshop on Logic Programming and Knowledge Representation (LPKR '97)*, October 17, 1997, Port Jefferson, New York, USA, volume 1471 of *Lecture Notes in Computer Science*, 224–246. Springer.
- Leite, J. A. 2003. *Evolving Knowledge Bases*, volume 81 of *Frontiers of Artificial Intelligence and Applications*, xviii + 307 p. Hardcover. IOS Press.
- Osorio, M., and Cuevas, V. 2007. Updates in answer set programming: An approach based on basic structural properties. *Theory and Practice of Logic Programming* 7(4):451–479.
- Sakama, C., and Inoue, K. 2003. An abductive framework for computing knowledge base updates. *Theory and Practice of Logic Programming (TPLP)* 3(6):671–713.
- Šefránek, J. 2006. Irrelevant updates and nonmonotonic assumptions. In Fisher, M.; van der Hoek, W.; Konev, B.; and Lisitsa, A., eds., *Proceedings of the 10th European Conference on Logics in Artificial Intelligence (JELIA 2006)*, volume 4160 of *Lecture Notes in Computer Science*, 426–438. Liverpool, UK: Springer.
- Šefránek, J. 2011. Static and dynamic semantics: Preliminary report. *Mexican International Conference on Artificial Intelligence* 36–42.
- Slota, M., and Leite, J. 2010. On semantic update operators for answer-set programs. In Coelho, H.; Studer, R.; and Wooldridge, M., eds., *ECAI 2010 - 19th European Conference on Artificial Intelligence, Lisbon, Portugal, August 16-20, 2010, Proceedings*, volume 215 of *Frontiers in Artificial Intelligence and Applications*, 957–962. IOS Press.
- Slota, M., and Leite, J. 2012a. Robust equivalence models for semantic updates of answer-set programs. In Brewka, G.; Eiter, T.; and McIlraith, S. A., eds., *Proceedings of the 13th International Conference on Principles of Knowledge Representation and Reasoning (KR 2012)*, 158–168. Rome, Italy: AAAI Press.
- Slota, M., and Leite, J. 2012b. A unifying perspective on knowledge updates. In del Cerro, L. F.; Herzig, A.; and Mengin, J., eds., *Logics in Artificial Intelligence - 13th European Conference, JELIA 2012, Toulouse, France, September 26-28, 2012. Proceedings*, volume 7519 of *Lecture Notes in Computer Science*, 372–384. Springer.
- Slota, M., and Leite, J. 2013. On condensing a sequence of updates in answer-set programming. In Rossi, F., ed., *IJCAI 2013, Proceedings of the 23rd International Joint Conference on Artificial Intelligence, Beijing, China, August 3-9, 2013*. IJCAI/AAAI.
- Slota, M., and Leite, J. 2014. The rise and fall of semantic rule updates based on se-models. *Theory and Practice of Logic Programming* FirstView:1–39.
- Slota, M.; Baláz, M.; and Leite, J. 2014. On strong and default negation in logic program updates (extended version). *CoRR* abs/1404.6784.
- Zhang, Y. 2006. Logic program-based updates. *ACM Transactions on Computational Logic* 7(3):421–472.