

When You Must Forget: beyond strong persistence when forgetting in answer set programming

RICARDO GONÇALVES, MATTHIAS KNORR, JOÃO LEITE

NOVA LINCS, Universidade Nova de Lisboa, Portugal

STEFAN WOLTRAN

TU Wien, Austria

submitted 2 May 2017; revised 20 June 2017; accepted 4 July 2017

Abstract

Among the myriad of desirable properties discussed in the context of forgetting in Answer Set Programming (ASP), *strong persistence* naturally captures its essence. Recently, it has been shown that it is not always possible to forget a set of atoms from a program while obeying this property, and a precise criterion regarding what can be forgotten has been presented, accompanied by a class of forgetting operators that return the correct result when forgetting is possible. However, it is an open question what to do when we have to forget a set of atoms, but cannot without violating this property. In this paper, we address this issue and investigate three natural alternatives to forget when forgetting without violating strong persistence is not possible, which turn out to correspond to the different possible relaxations of the characterization of strong persistence. Additionally, we discuss their preferable usage, shed light on the relation between forgetting and notions of relativized equivalence established earlier in the context of ASP, and present a detailed study on their computational complexity.

KEYWORDS: Forgetting, Answer Set Programming, Strong Equivalence, Relativized Equivalence, Computational Complexity

1 Introduction

A fundamental conclusion drawn by Gonçalves et al. (2016c) is that it is sometimes impossible to *forget* a set of atoms from an answer set program while obeying important desirable properties, notably the so-called *strong persistence*. However, even in such cases, we may be forced to forget – just imagine a court ordering the elimination of illegally acquired information. In this paper, we thoroughly investigate how to forget when it is impossible to obey strong persistence.

Forgetting is an operation that allows the removal from a knowledge base of *middle* variables no longer deemed relevant. Its importance is witnessed by its application, e.g., to cognitive robotics (Lin and Reiter 1997; Liu and Wen 2011; Rajaratnam et al. 2014), resolving conflicts (Lang et al. 2003; Zhang and Foo 2006; Eiter and Wang 2008; Lang and Marquis 2010), and ontology abstraction and comparison (Wang et al. 2010; Kontchakov et al. 2010; Konev et al. 2012; Konev et al. 2013). With its early roots in Boolean Algebra (Lewis 1918), it has been extensively studied within classical logic (Bledsoe and Hines 1980; Lang et al. 2003; Larrosa 2000; Larrosa et al. 2005; Middeldorp et al. 1996; Moïnard 2007; Weber 1986) and, more recently, in

the context of Answer Set Programming (ASP) (Leite 2017). The non-monotonic rule-based nature of ASP creates very unique challenges to the development of forgetting operators – just as it happened with other belief change operations such as revision and update, cf. (Alferes et al. 2000; Eiter et al. 2002; Sakama and Inoue 2003; Slota and Leite 2012; Delgrande et al. 2013; Slota and Leite 2014) – making it a special endeavour with unique characteristics distinct from those for classical logic. This led to the introduction of several forgetting operators and classes of operators (Zhang and Foo 2006; Eiter and Wang 2008; Wong 2009; Wang et al. 2012; Wang et al. 2013; Knorr and Alferes 2014; Wang et al. 2014; Delgrande and Wang 2015; Gonçalves et al. 2016c), often promoting different sets of properties that such operators of forgetting should satisfy.

From a recent critical survey on forgetting in ASP (Gonçalves et al. 2016b), it stood out that *strong persistence* (**SP**) (Knorr and Alferes 2014) – a property essentially requiring that all existing relations between the atoms not to be forgotten be preserved – best captures the essence of forgetting in the context of ASP. However, as shown by Gonçalves et al. (2016c), sometimes the atoms to be forgotten play such a pivotal role that they cannot be forgotten without violating (**SP**). The class of situations when forgetting is possible was characterized through a criterion – Ω – that can be applied to any answer set program P and set of atoms V , holding whenever V cannot be forgotten from P , and not holding otherwise. For those cases when forgetting is possible, Gonçalves et al. (2016c) also present a class of forgetting operators that satisfy (**SP**), dubbed F_{SP} .

But what if Ω is true and we nevertheless *must* forget? This may happen for legal and privacy issues, including, for example, to enforce the new EU General Data Protection Regulation (European Parliament 2016), which includes the *right to be forgotten* – the person’s right to ask a corporation to eliminate private data – or the implementation of court orders to eliminate certain pieces of illegally acquired or maintained information. Tools that can help companies and users automate the operation of forgetting should be able to handle not only situations where we can achieve the required forgetting without violating *strong persistence*, but also situations where such *ideal* forgetting is not possible. Towards developing a theoretical ground on which such universally applicable tools can be based, in this paper, we thoroughly address the question of how to forget when Ω is true, along three different ways.

We *first* take a closer look at the class F_{SP} , which had only been considered for the case when Ω is false, and investigate how it behaves in general. One crucial observation is that it overestimates answer sets, i.e., forgetting preserves all existing answer sets, but new ones may be added, which indicates a violation of property (**sC**) (strengthened consequence).

Our *second* approach borrows from the notion of *relativized equivalence* (Eiter et al. 2007), a generalization of strong equivalence that considers equivalence only w.r.t. a given subset of the language, and is characterized by the so-called V -HT-models¹, which lead us to consider two novel ways to forget: a specific operator that simply returns all rules that are relativized equivalent to the original program w.r.t. the atoms not to be forgotten and, alternatively, a class of operators whose result is characterized by the set of V -HT-models, omitting the atoms to be forgotten. The former operator turns out to be a member of the latter class. Whereas this class never overestimates answer sets, i.e., it obeys (**sC**), it may lose some of the original answer sets, which indicates a violation of property (**wC**) (weakened consequence).

¹ Programs P_1, P_2 are *relativized equivalent* w.r.t. $V \subseteq \mathcal{A}$ if and only if they have the same V -HT-models

The *third* approach tries to overcome a weakness of the second, i.e., its result diverges from F_{SP} even when it is possible to forget, and proposes a case-based definition that can be seen as a mixture of the previous two. Whereas it preserves all answer sets, i.e., it obeys both **(sC)** and **(wC)**, it no longer satisfies **(SI)** (strong invariance), i.e., forgetting first and then adding some set of rules R (not containing the atoms to be forgotten) is no longer (strongly) equivalent to adding R first, and forgetting subsequently.

In this paper, we fully investigate these three alternatives. We characterize them by showing which subset of the properties previously considered in the literature each of them obeys, study their computational complexity, and relate them by considering further additional properties to help clarify their preferable usage. Perhaps one of the most interesting features of this set of alternatives stems from a characterisation of **(SP)** according to which a forgetting operator obeys **(SP)** if and only if it obeys **(sC)**, **(wC)** and **(SI)**. Hence, each of the three alternatives exactly corresponds to the relaxation of one of these three properties that characterize **(SP)**.

Additional relevant results include a formal correspondence between V -HT-models and HT-models allowing us to leverage beneficial properties of HT-models, such as monotonicity, in the realm of V -HT-models, which do not satisfy them, and a complexity result for checking whether Ω holds.

The remainder of the paper starts with some background on forgetting in ASP, then proceeds with one section for each of the three approaches, followed by one on their complexity, and one with some brief concluding remarks.

2 Forgetting in ASP

In this section, we recall the necessary notions on answer set programming and forgetting.

Logic programs We assume a *propositional signature* \mathcal{A} , a finite set of propositional atoms². An (*extended*) *logic program* P over \mathcal{A} is a finite set of (*extended*) *rules* of the form

$$a_1 \vee \dots \vee a_k \leftarrow b_1, \dots, b_l, \text{not } c_1, \dots, \text{not } c_m, \text{not not } d_1, \dots, \text{not not } d_n, \quad (1)$$

where all $a_1, \dots, a_k, b_1, \dots, b_l, c_1, \dots, c_m$, and d_1, \dots, d_n are atoms of \mathcal{A} .³ Such rules r are also commonly written in a more succinct way as

$$A \leftarrow B, \text{not } C, \text{not not } D, \quad (2)$$

where we have $A = \{a_1, \dots, a_k\}$, $B = \{b_1, \dots, b_l\}$, $C = \{c_1, \dots, c_m\}$, $D = \{d_1, \dots, d_n\}$, and we will use both forms interchangeably. By $\mathcal{A}(P)$ we denote the set of atoms appearing in P . This class of logic programs, \mathcal{C}_e , includes a number of special kinds of rules r : if $n = 0$, then we call r *disjunctive*; if, in addition, $k \leq 1$, then r is *normal*; if on top of that $m = 0$, then we call r *Horn*, and *fact* if also $l = 0$. The classes of *disjunctive*, *normal* and *Horn programs*, \mathcal{C}_d , \mathcal{C}_n , and \mathcal{C}_H , are defined resp. as a finite set of disjunctive, normal, and Horn rules. Given a program P and a set I of atoms, the *reduct* P^I is defined as $P^I = \{A \leftarrow B : r \text{ of the form (2) in } P, C \cap I = \emptyset, D \subseteq I\}$.

An *HT-interpretation* is a pair $\langle X, Y \rangle$ s.t. $X \subseteq Y \subseteq \mathcal{A}$. Given a program P , an HT-interpretation $\langle X, Y \rangle$ is an *HT-model* of P if $Y \models P$ and $X \models P^Y$, where \models denotes the standard consequence relation for classical logic. We admit that the set of HT-models of a program P are restricted to

² Often, the term propositional variable is used synonymously.

³ Extended logic programs (Lifschitz et al. 1999) are actually more expressive, but this form is sufficient here.

$\mathcal{A}(P)$ even if $\mathcal{A}(P) \subset \mathcal{A}$. We denote by $\mathcal{HT}(P)$ the set of all *HT-models* of P . A set of atoms Y is an *answer set* of P if $\langle Y, Y \rangle \in \mathcal{HT}(P)$, but there is no $X \subset Y$ such that $\langle X, Y \rangle \in \mathcal{HT}(P)$. The set of all answer sets of P is denoted by $\mathcal{AS}(P)$. We say that two programs P_1, P_2 are *equivalent* if $\mathcal{AS}(P_1) = \mathcal{AS}(P_2)$ and *strongly equivalent*, denoted by $P_1 \equiv P_2$, if $\mathcal{AS}(P_1 \cup R) = \mathcal{AS}(P_2 \cup R)$ for any $R \in \mathcal{C}_e$. It is well-known that $P_1 \equiv P_2$ exactly when $\mathcal{HT}(P_1) = \mathcal{HT}(P_2)$ (Lifschitz et al. 2001). We say that P' is an *HT-consequence* of P , denoted by $P \models_{\text{HT}} P'$, whenever $\mathcal{HT}(P) \subseteq \mathcal{HT}(P')$. The *V-exclusion* of a set of answer sets (a set of HT-interpretations) \mathcal{M} , denoted $\mathcal{M}_{\parallel V}$, is $\{X \setminus V \mid X \in \mathcal{M}\} \cup \{\langle X \setminus V, Y \setminus V \rangle \mid \langle X, Y \rangle \in \mathcal{M}\}$. Finally, given two sets of atoms $X, X' \subseteq \mathcal{A}$, we write $X \sim_V X'$ whenever $X \setminus V = X' \setminus V$.

We recall the notion of *A-SE-models* (Eiter et al. 2007), but here adapted to *V-HT-models* that focus on $V \subseteq \mathcal{A}$, instead of on $A = \mathcal{A} \setminus V$. An HT-interpretation $\langle X, Y \rangle$ is called a *V-HT-interpretation* if either $X = Y$ or $X \subset Y \setminus V$. A *V-HT-interpretation* $\langle X, Y \rangle$ is a (*relativized*) *V-HT-model* of P if: (a) $Y \models P$; (b) for all $Y' \subset Y$ with $Y \sim_V Y'$, $Y' \not\models P^Y$; and (c) if $X \subset Y$, then there exists $X' \subseteq Y$ such that $X = X' \setminus V$ and $X' \models P^Y$. We denote by $\mathcal{HT}_V(P)$ the set of all *V-HT-models* of P . Programs P_1, P_2 are *relativized equivalent* w.r.t. $V \subseteq \mathcal{A}$, denoted by $P_1 \equiv_V P_2$, if $\mathcal{AS}(P_1 \cup R) = \mathcal{AS}(P_2 \cup R)$ for any $R \in \mathcal{C}_e$ s.t. $\mathcal{A}(R) \subseteq \mathcal{A} \setminus V$. We have that $P_1 \equiv_V P_2$ exactly when $\mathcal{HT}_V(P_1) = \mathcal{HT}_V(P_2)$ (Eiter et al. 2007).

Forgetting Given a class of logic programs \mathcal{C} over \mathcal{A} , a *forgetting operator* (over \mathcal{C}) is a partial function $f : \mathcal{C} \times 2^{\mathcal{A}} \rightarrow \mathcal{C}$ s.t. $f(P, V)$ is a program over $\mathcal{A}(P) \setminus V$, for each $P \in \mathcal{C}$ and $V \subseteq \mathcal{A}$. We call $f(P, V)$ the *result of forgetting about V from P*. Unless stated otherwise, in what follows, we will be focusing on $\mathcal{C} = \mathcal{C}_e$, and we leave \mathcal{C} implicit. Furthermore, f is called *closed* for $\mathcal{C}' \subseteq \mathcal{C}$ if, for every $P \in \mathcal{C}'$ and $V \subseteq \mathcal{A}$, we have $f(P, V) \in \mathcal{C}'$. A *class F of forgetting operators* (over \mathcal{C}) is a set of forgetting operators (over \mathcal{C}') s.t. $\mathcal{C}' \subseteq \mathcal{C}$. Such classes are usually described by a common definition/condition that each operator in the class has to satisfy (see the paper by Gonçalves et al. (2016b) for an overview on the many different kinds and forms of defining such classes).

At the same time, previous work on forgetting in ASP has introduced a variety of desirable properties accompanying these classes of operators. In the following, we recall these properties and leave the details, e.g., on which class of forgetting operators satisfies which properties to the papers by Gonçalves et al. (2016b; 2016a).⁴ Unless stated otherwise, F is a class of forgetting operators, and \mathcal{C} the class of programs over \mathcal{A} of a given $f \in F$.

- (sC) F satisfies *strengthened Consequence* if, for each $f \in F$, $P \in \mathcal{C}$ and $V \subseteq \mathcal{A}$, we have $\mathcal{AS}(f(P, V)) \subseteq \mathcal{AS}(P)_{\parallel V}$.
- (wE) F satisfies *weak Equivalence* if, for each $f \in F$, $P, P' \in \mathcal{C}$ and $V \subseteq \mathcal{A}$, we have $\mathcal{AS}(f(P, V)) = \mathcal{AS}(f(P', V))$ whenever $\mathcal{AS}(P) = \mathcal{AS}(P')$.
- (SE) F satisfies *Strong Equivalence* if, for each $f \in F$, $P, P' \in \mathcal{C}$ and $V \subseteq \mathcal{A}$: if $P \equiv P'$, then $f(P, V) \equiv f(P', V)$.
- (W) F satisfies *Weakening* if, for each $f \in F$, $P \in \mathcal{C}$ and $V \subseteq \mathcal{A}$, we have $P \models_{\text{HT}} f(P, V)$.
- (PP) F satisfies *Positive Persistence* if, for each $f \in F$, $P \in \mathcal{C}$ and $V \subseteq \mathcal{A}$: if $P \models_{\text{HT}} P'$, with $P' \in \mathcal{C}$ and $\mathcal{A}(P') \subseteq \mathcal{A} \setminus V$, then $f(P, V) \models_{\text{HT}} P'$.
- (SI) F satisfies *Strong (addition) Invariance* if, for each $f \in F$, $P \in \mathcal{C}$ and $V \subseteq \mathcal{A}$, we have $f(P, V) \cup R \equiv f(P \cup R, V)$ for all programs $R \in \mathcal{C}$ with $\mathcal{A}(R) \subseteq \mathcal{A} \setminus V$.

⁴ We omit (NP) from the list, as it has been shown there to coincide with (W).

- (**E_C**) F satisfies *Existence for C*, i.e., F is closed for a class of programs \mathcal{C} if there exists $f \in F$ s.t. f is closed for \mathcal{C} .
- (**CP**) F satisfies *Consequence Persistence* if, for each $f \in F$, $P \in \mathcal{C}$ and $V \subseteq \mathcal{A}$, we have $\mathcal{AS}(P)_{\parallel V} = \mathcal{AS}(f(P, V))$.
- (**wC**) F satisfies *weakened Consequence* if, for each $f \in F$, $P \in \mathcal{C}$ and $V \subseteq \mathcal{A}$, we have $\mathcal{AS}(P)_{\parallel V} \subseteq \mathcal{AS}(f(P, V))$.
- (**SP**) F satisfies *Strong Persistence* if, for each $f \in F$, $P \in \mathcal{C}$ and $V \subseteq \mathcal{A}$, we have $\mathcal{AS}(f(P, V) \cup R) = \mathcal{AS}(P \cup R)_{\parallel V}$, for all programs $R \in \mathcal{C}$ with $\mathcal{A}(R) \subseteq \mathcal{A} \setminus V$.

We refer to the recent critical survey by Gonçalves et al. (2016b) for a discussion about existing relations between these properties, but we want to point out that the importance of (**SP**) is witnessed by the fact that if some class F satisfies (**SP**), then it also satisfies basically all other mentioned properties (but (**W**) and (**E_C**), which is orthogonal).

Example 1

Consider the following program P .

$$a \leftarrow \text{not } b \qquad b \leftarrow \text{not } c \qquad e \leftarrow d \qquad d \leftarrow a$$

First, if we want to forget about an atom, then we expect that all rules that do not mention this atom should persist, while rules that do mention the atoms should no longer occur. For example, when forgetting about d from P , the first two rules should be contained in the result of the forgetting, while the latter two should not. At the same time, implicit dependencies should be preserved, such as, e depending on a via d . Hence, we expect $f(P, \{d\})$ as follows:

$$a \leftarrow \text{not } b \qquad b \leftarrow \text{not } c \qquad e \leftarrow a$$

In fact, many existing notions of forgetting in the literature (cf. the paper by Gonçalves et al. (2016b)) provide precisely this result.

Now, consider forgetting about b from P . Note that P contains an implicit dependency between a and c , namely, whenever c becomes true, then so does a , i.e., if we add, e.g., $c \leftarrow$ to the program, then a is necessarily true. Different notions of (classes of) forgetting operators f existing in the literature (see the paper by Gonçalves et al. (2016b)) would return the result of forgetting $f(P, \{b\}) = \emptyset$, but if we want to preserve property (**SP**), then $f(P, \{b\})$ must contain the rule $a \leftarrow \text{not } c$. In fact, a valid result for $f(P, \{b\})$ such that f satisfies (**SP**) is:

$$a \leftarrow \text{not } c \qquad e \leftarrow d \qquad d \leftarrow a$$

Finally, if the atom to be forgotten does not appear at the same time in some rule body and some rule head, usually no dependencies need to be preserved. Consider forgetting about c from P , then, since c only appears in the body of a rule, the result $f(P, \{c\})$ is:

$$a \leftarrow \text{not } b \qquad b \leftarrow \qquad e \leftarrow d \qquad d \leftarrow a$$

3 On the Limits of Forgetting

As argued by Gonçalves et al. (2016c), (**SP**) is the central property one wants to ensure to hold when forgetting atoms from an answer set program, essentially because its definition intuitively requires that all (direct and indirect) dependencies between the atoms not to be forgotten be preserved. This is witnessed by the fact that any class of forgetting operators that satisfies (**SP**) also satisfies all other properties introduced in the literature, with the exception of (**W**), which

has been shown to be incompatible with **(SP)** (Gonçalves et al. 2016b). However, it is also shown that it is not always possible to forget a set of atoms from a given program, that is, there is no forgetting operator that satisfies **(SP)** and that is defined for all pairs $\langle P, V \rangle$, called *forgetting instances*, where P is a program and V is a set of atoms to be forgotten from P . The precise characterization of when it is not possible to forget while satisfying **(SP)** is given by means of criterion Ω .

Definition 1 (Criterion Ω)

Let P be a program over \mathcal{A} and $V \subseteq \mathcal{A}$. An instance $\langle P, V \rangle$ satisfies criterion Ω if there exists $Y \subseteq \mathcal{A} \setminus V$ such that the set of sets

$$\mathcal{R}_{\langle P, V \rangle}^Y = \{R_{\langle P, V \rangle}^{Y, A} \mid A \in \text{Rel}_{\langle P, V \rangle}^Y\}$$

is non-empty and has no least element, where

$$\begin{aligned} R_{\langle P, V \rangle}^{Y, A} &= \{X \setminus V \mid \langle X, Y \cup A \rangle \in \mathcal{HT}(P)\} \\ \text{Rel}_{\langle P, V \rangle}^Y &= \{A \subseteq V \mid \langle Y \cup A, Y \cup A \rangle \in \mathcal{HT}(P) \text{ and} \\ &\quad \nexists A' \subset A \text{ s.t. } \langle Y \cup A', Y \cup A \rangle \in \mathcal{HT}(P)\}. \end{aligned}$$

The rationale is that each set $\mathcal{R}_{\langle P, V \rangle}^Y$ is based on $Y \subseteq \mathcal{A} \setminus V$, which is a potential answer set of the result of forgetting. Taking property **(SP)** into account, an answer set Y of $f(P, V) \cup R$ must be obtained from an answer set $Y \cup A$ of $P \cup R$, for some $A \subseteq V$. So, the HT-models of the form $\langle X, Y \cup A \rangle$ in $\mathcal{HT}(P)$ must be taken into account. This is captured by the set $\mathcal{R}_{\langle P, V \rangle}^{Y, A}$. Nevertheless, there are some $A \subseteq V$ such that $Y \cup A$ is never an answer set of $P \cup R$, for any R over $\mathcal{A} \setminus V$. This is captured by the condition of the set $\text{Rel}_{\langle P, V \rangle}^Y$.

This criterion was shown to be sound and complete, i.e., it is not possible to forget about a set of atoms V from a program P exactly when $\langle P, V \rangle$ satisfies criterion Ω . A corresponding class of forgetting operators, F_{SP} , was introduced.

Definition 2 (SP-Forgetting)

Let F_{SP} be the class of forgetting operators defined by the following set:⁵

$$\{f \mid \mathcal{HT}(f(P, V)) = \{\langle X, Y \rangle \mid Y \subseteq \mathcal{A} \setminus V \wedge X \in \bigcap \mathcal{R}_{\langle P, V \rangle}^Y\}\}$$

It was shown that every operator in F_{SP} satisfies **(SP)** for instances that do not satisfy Ω . In fact, restricted to those instances, F_{SP} satisfies every property except **(W)**, which makes this class of operators an ideal choice whenever forgetting is possible.

However, the question as to whether this class is also of any use in case Ω is satisfied has not been tackled. Given our focus on this problem, we first consider F_{SP} itself as a possible solution and characterize which of the well-known properties of forgetting are satisfied by F_{SP} in general, i.e., independently of whether Ω is satisfied or not.

Proposition 1

F_{SP} satisfies **(wC)**, **(SE)**, **(PP)**, and **(SI)**, but does not satisfy **(wE)**, **(W)**, **(sC)**, **(CP)**.

Regarding *existence*, it has already been shown by Gonçalves et al. (2016c) that F_{SP} is closed for extended programs and Horn programs, but not for disjunctive nor normal programs.

⁵ The definition is slightly generalized from the paper by Gonçalves et al. (2016c) as Y is no longer restricted to be $Y \subseteq \mathcal{A}(P) \setminus V$. Whenever $\mathcal{A} = \mathcal{A}(P)$, then the two versions naturally coincide.

From the previous proposition, we observe that **(W)** is no longer the only property that does not hold. Notably, the fact that F_{SP} does not satisfy **(CP)**, and in particular **(sC)**, means that there are instances $\langle P, V \rangle$ for which the result of forgetting about V from P has answer sets that do not correspond to answer sets in the original program P , which is also why **(wE)** does not hold.

Example 2

Consider the following program P .

$$a \leftarrow p \qquad b \leftarrow \text{not } p \qquad p \leftarrow \text{not not } p$$

Clearly, P has six HT-models, $\langle ap, ap \rangle, \langle b, b \rangle, \langle b, ab \rangle, \langle ab, ab \rangle, \langle ap, abp \rangle, \langle abp, abp \rangle$ ⁶, and two answer sets $\{a, p\}$ and $\{b\}$. Intuitively, p yields an exclusive choice between a and b . If we take $V = \{p\}$, then, $\mathcal{R}_{\langle P, V \rangle}^{\emptyset} = \emptyset$, $\mathcal{R}_{\langle P, V \rangle}^{\{a\}} = \{\{a\}\}$, $\mathcal{R}_{\langle P, V \rangle}^{\{b\}} = \{\{b\}\}$, and $\mathcal{R}_{\langle P, V \rangle}^{\{a, b\}} = \{\{b, ab\}, \{a, ab\}\}$. From this we have that $\bigcap \mathcal{R}_{\langle P, V \rangle}^{\emptyset} = \emptyset$, $\bigcap \mathcal{R}_{\langle P, V \rangle}^{\{a\}} = \{a\}$, $\bigcap \mathcal{R}_{\langle P, V \rangle}^{\{b\}} = \{b\}$, and $\bigcap \mathcal{R}_{\langle P, V \rangle}^{\{a, b\}} = \{ab\}$. This means that for any $f \in F_{SP}$, $f(P, V)$ has three HT-models, $\langle a, a \rangle, \langle b, b \rangle, \langle ab, ab \rangle$, which means that $f(P, V)$ has three answer sets, the two from P ignoring p , $\{a\}$ and $\{b\}$, and additionally $\{a, b\}$. Intuitively, this happens because using the intersection essentially discards both $\langle b, ab \rangle$ and $\langle ap, abp \rangle$ (modulo the forgotten p).

This is in fact rather atypical as so far no class of forgetting operators that satisfies **(wC)**, but not **(sC)**, and thus not **(CP)**, was known. Since the violation of **(sC)** may be seen as sufficient cause to render F_{SP} inadequate when Ω is satisfied – notably when the introduction of new answer sets as the result of forgetting cannot be accepted – alternatives need to be investigated.

4 Relativized Forgetting

In this section, we explore alternative ways to forgetting in ASP, borrowing from the notion of relativized equivalence (Eiter et al. 2007). Relativized equivalence is a generalization of strong equivalence that considers equivalence w.r.t. a given subset of the language, such that equivalence and strong equivalence are its special cases (for the empty and the entire language respectively). This fits naturally within the idea of forgetting in ASP, in particular w.r.t. property **(SP)**, inasmuch as after forgetting about V from P we only allow the addition of programs over $\mathcal{A} \setminus V$, so relativized (strong) equivalence should be applied accordingly.

Based on this idea, we first define a forgetting operator that simply considers all logical consequences w.r.t. relativized equivalence. This way, the result of forgetting about V from P amounts to the set of all rules (over $\mathcal{A} \setminus V$) that can be added to P while preserving relativized equivalence. Given a program P and $V \subseteq \mathcal{A}$, we consider the closure of P given V :

$$Cn(P, V) = \{r \mid \{r\} \in \mathcal{C}_e \text{ and } P \cup \{r\} \equiv_V P\}.$$

Then, the result of forgetting about V from P is defined as

$$f_r(P, V) = \{r \mid r \in Cn(P, V) \text{ and } \mathcal{A}(\{r\}) \cap V = \emptyset\}.$$

The resulting program does not mention the forgotten atoms and we can show that this operator does not belong to F_{SP} .

⁶ We follow a common convention and abbreviate sets in HT-interpretations such as $\{a, b\}$ with the sequence of its elements, ab .

Example 3

Recall program P from Ex. 2. It can be verified that $f_r(P, \{p\})$ is strongly equivalent to the program:

$$\begin{array}{ll} a \leftarrow \text{not } b & \perp \leftarrow \text{not } a, \text{not } b \\ b \leftarrow \text{not } a & a \vee b \leftarrow \end{array}$$

Notably, this program does not have the answer set $\{a, b\}$, which indicates that this operator does not belong to F_{SP} .

We can show that f_r is well-defined, in the sense that testing relativized equivalence for each rule individually is the same as testing the entire set of rules as a whole.

Proposition 2

Let P be a program, $V \subseteq \mathcal{A}$ and R_1, R_2 programs over $\mathcal{A} \setminus V$. Then, $P \cup R_1 \cup R_2 \equiv_V P$ iff $P \cup R_1 \equiv_V P$ and $P \cup R_2 \equiv_V P$.

As a consequence of the above result, $f_r(P, V)$ is in fact the largest set of rules over $\mathcal{A} \setminus V$ that can be safely added to P without changing its set of V -HT-models.

Proposition 3

Let P be a program and $V \subseteq \mathcal{A}$. Then, $f_r(P, V)$ is the largest set of rules R over the alphabet $\mathcal{A} \setminus V$ such that $P \cup R \equiv_V P$.

We could now define a (possibly singleton) class of operators that generalizes the idea of f_r in a straightforward manner, and then study this class, but its definition would not be very concise, as we would always have to check for each rule whether it is relativized equivalent to the original program.

Instead, inspired by knowledge forgetting (Wang et al. 2014), we follow a different idea, defining a class of forgetting operators that consider the V -HT-models of P and omit all occurrences of elements of V from these. Formally:

$$F_R = \{f \mid \mathcal{HT}(f(P, V)) = \mathcal{HT}_V(P)_{\parallel V}\}$$

Example 4

Recall Ex. 3. It can be verified that the result of forgetting for any $f \in F_R$ coincides with that for f_r .

It turns out that this correspondence is no mere coincidence. In fact, we show in the following that $f_r \in F_R$, and in the course of that, we establish a precise relation between the HT-models and the V -HT-models of a program. This is an important contribution, since it allows the usage of well-known properties of HT-models, such as monotonicity, that are not satisfied by V -HT-models (see the paper by Eiter et al. (2007)).

First, we introduce an alternative characterization of the V -HT-models of a program P based on its HT-models using the following notion.

Definition 3

Let P be a program and $Y, V \subseteq \mathcal{A}$. Then, Y is *relevant for P w.r.t. V* if

- (i) $\langle Y, Y \rangle \in \mathcal{HT}(P)$
- (ii) $\langle Y', Y \rangle \notin \mathcal{HT}(P)$ for every $Y' \subset Y$ s.t. $Y \sim_V Y'$.

$Rel(P, V)$ denotes the set of all sets relevant for P w.r.t. V .

This notion is tightly connected with the sets in the definition of criterion Ω , i.e., we can show that $Y \cup A \in \text{Rel}(P, V)$ iff $A \in \text{Rel}_{\langle P, V \rangle}^Y$. This allows the alternative definition of a V -HT-model in terms of HT-models.

Proposition 4

Let P be a program and $V \subseteq \mathcal{A}$. Then, a V -HT-interpretation $\langle X, Y \rangle$ is a V -HT-model of P iff the following conditions hold:

- (1) $Y \in \text{Rel}(P, V)$;
- (2) If $X \subset Y$, then there exists $X' \subset Y$ with $X = X' \setminus V$ such that $\langle X', Y \rangle \in \mathcal{HT}(P)$.

We can now present an alternative characterization of the set of V -HT-models of a program in terms of its set of HT-models. This result is particularly useful since it shows how the set of V -HT-models of a program can be directly obtained from its set of HT-models.

Proposition 5

Let P be a program and $V \subseteq \mathcal{A}$. Then,

$$\mathcal{HT}_V(P) = \bigcup_{Y \in \text{Rel}(P, V)} (\{\langle X \setminus V, Y \rangle : \langle X, Y \rangle \in \mathcal{HT}(P) \text{ and } X \subset Y\} \cup \{\langle Y, Y \rangle\}).$$

Based on that, we can show that f_r is indeed a concrete forgetting operator in the class F_R .

Theorem 1

Let P be a program and $V \subseteq \mathcal{A}$. Then,

$$\mathcal{HT}(f_r(P, V))_{\parallel V} = \mathcal{HT}_V(P)_{\parallel V}.$$

Interestingly, we are also able to provide an alternative characterization of F_R that clarifies the relation to F_{SP} .

Theorem 2

Let P be a program and $V \subseteq \mathcal{A}$. Then, F_R can be given by the set

$$\{f \mid \mathcal{HT}(f(P, V)) = \{\langle X, Y \rangle \mid Y \subseteq \mathcal{A} \setminus V \wedge X \in \bigcup \mathcal{R}_{\langle P, V \rangle}^Y\}\}.$$

Thus, this notion of forgetting based on relativized equivalence differs from F_{SP} by considering the union of the relevant HT-models instead of the intersection, which explains the differences observed in Ex. 2 and 3.

Of course, whenever $\mathcal{R}_{\langle P, V \rangle}^Y$ contains only one element, union and intersection coincide, which is always the case for Horn programs.

Proposition 6

Let $P \in \mathcal{C}_H$ and $V \subseteq \mathcal{A}$. Then, for every $Y \subseteq \mathcal{A} \setminus V$, we have that $\mathcal{R}_{\langle P, V \rangle}^Y$ has at most one element.

Thus, when restricted to \mathcal{C}_H , F_R coincides with F_{SP} .

Proposition 7

Let $P \in \mathcal{C}_H$ and $V \subseteq \mathcal{A}$. Then, for every $f \in F_{SP}$ and $f' \in F_R$ we have that $f(P, V) \equiv f'(P, V)$.

Since this correspondence does not hold in general, we also establish which properties are satisfied by F_R .

Proposition 8

F_R satisfies **(sC)**, **(SE)**, **(PP)**, **(SI)**, **(E_{CH})**, **(E_{C_e)}**, but not **(wE)**, **(W)**, **(wC)**, **(CP)**, **(E_{C_n)}**, **(E_{C_d)}**.

In terms of the considered set of properties, F_R and F_{SP} only differ with respect to **(sC)** and **(wC)**. This difference, however, is crucial. Since F_R satisfies **(sC)**, it approximates the set of answer sets of P , but, contrary to F_{SP} , never ends up adding new answer sets to the result of forgetting. However, it's not all roses, as will become clear next.

5 Merging F_{SP} and F_R

We have shown that F_R , which is based on relativized forgetting, is a better alternative than F_{SP} if our objective is to approximate the set of answer sets modulo the forgotten atoms, but not introduce new answer sets. However, F_R has a drawback: there are cases where it is possible to forget while satisfying **(SP)**, but the result for any $f \in F_R$ does not coincide with the desired result (obtainable with operators from F_{SP}).

Example 5

Consider the following program P and that we want to forget about p from P .

$$a \leftarrow p \qquad p \leftarrow \text{not not } p$$

It is easy to check that $\langle P, V \rangle$ does not satisfy Ω , i.e., it is possible to forget about V from P while satisfying **(SP)**. The result returned by any operator in F_{SP} is strongly equivalent to $\{a \leftarrow \text{not not } a\}$. However, $f(P, V)$ for any $f \in F_R$ is strongly equivalent to the empty program.

The difference between F_{SP} and F_R , as shown in Thm. 2, lies in the usage of intersection and union in their respective definitions. The key point is that whenever $\mathcal{R}_{\langle P, V \rangle}^Y$ has more than one element, even if there is a least one, union and intersection will not coincide. Taking this idea into account, we define a class of operators that aims at combining the delineated positive aspects of both F_{SP} and F_R .

$$F_M = \{f \mid \mathcal{HT}(f(P, V)) = \{\langle X, Y \rangle \mid Y \subseteq \mathcal{A} \setminus V \text{ and} \\ X \in \bigcup \mathcal{R}_{\langle P, V \rangle}^Y, \text{ if } \mathcal{R}_{\langle P, V \rangle}^Y \text{ has no least element, or} \\ X \in \bigcap \mathcal{R}_{\langle P, V \rangle}^Y, \text{ otherwise}\}\}.$$

Whenever $\mathcal{R}_{\langle P, V \rangle}^Y$ has a least element, then F_M employs the intersection, whose result is precisely the least element, similar to F_{SP} and does therefore coincide with the desired ideal solution in this case, and whenever there is no least element it uses the union instead, just like F_R .

Example 6

Consider the program of Ex. 2. The result of forgetting about p from that program, for any $f \in F_M$, is strongly equivalent with that given in Ex. 3 for any $f' \in F_R$. On the other hand, for the program given in Ex. 5, the result of forgetting about p from that program, for any $f \in F_M$, is strongly equivalent to $\{a \leftarrow \text{not not } a\}$, and the same also holds for any operator in F_{SP} .

Still, if we consider only Horn programs, then this definition of F_M coincides with both its constituents.

Proposition 9

Let $P \in \mathcal{C}_H$ and $V \subseteq \mathcal{A}$. Then, for every $f \in (F_{SP} \cup F_R)$ and $f' \in F_M$ we have that $f(P, V) \equiv f'(P, V)$.

Moreover, unlike F_R , we are able to show that, whenever it is possible to forget, F_M coincides with F_{SP} .

Proposition 10

Let P be a program and $V \subseteq \mathcal{A}$, such that $\langle P, V \rangle$ does not satisfy Ω . Then, for every $f \in F_{SP}$ and $f' \in F_M$ we have that $f(P, V) \equiv f'(P, V)$.

The particular definition of F_M ensures that yet again a different set of properties is satisfied by it.

Proposition 11

F_M satisfies **(sC)**, **(wE)**, **(SE)**, **(wC)**, **(CP)**, **(PP)**, **(E_{C_H)}**, **(E_{C_e)}**, but not **(W)**, **(SI)**, **(E_{C_n)}**, **(E_{C_d)}**.

Contrary to F_{SP} and F_R , the class F_M satisfies both **(wC)** and **(sC)**, and consequently **(CP)**. Therefore, the result of forgetting according to F_M preserves the answer sets of P , but, unlike the other two, no longer satisfies **(SI)**.

In fact, the answer sets are no longer preserved if a (non-empty) program over $\mathcal{A} \setminus V$ is added to P . To capture this in a more precise way, we introduce generalizations of **(wC)** and **(sC)**, which correspond to the two inclusions of **(SP)**.

(sSP) F satisfies *strengthened Strong Persistence* if, for each $f \in F$, $P \in \mathcal{C}$ and $V \subseteq \mathcal{A}$, we have $\mathcal{AS}(f(P, V) \cup R) \subseteq \mathcal{AS}(P \cup R)_{\parallel V}$, for all $R \in \mathcal{C}$ with $\mathcal{A}(R) \subseteq \mathcal{A} \setminus V$.

(wSP) F satisfies *weakened Strong Persistence* if, for each $f \in F$, $P \in \mathcal{C}$ and $V \subseteq \mathcal{A}$, we have $\mathcal{AS}(P \cup R)_{\parallel V} \subseteq \mathcal{AS}(f(P, V) \cup R)$, for all $R \in \mathcal{C}$ with $\mathcal{A}(R) \subseteq \mathcal{A} \setminus V$.

Property **(wSP)** guarantees that all answer sets of P are preserved when forgetting, no matter which rules R over $\mathcal{A} \setminus V$ are added to P , but, for some such R , does not prevent that the result of forgetting has more answer sets than P . Vice versa, **(sSP)** does not guarantee the preservation of all answer sets of P for some added R over $\mathcal{A} \setminus V$, but it ensures that all answer sets of the result of forgetting indeed correspond to answer sets of P , independently of the added rules R .

We can show that each of the three considered classes of forgetting operators only satisfies one of the two properties.

Theorem 3

F_{SP} satisfies **(wSP)**, whereas F_R and F_M satisfy **(sSP)**.

Since there is no class of forgetting operators that satisfies **(SP)** (Gonçalves et al. 2016c), it is clear that F_{SP} does not satisfy **(sSP)**, and that F_R and F_M do not satisfy **(wSP)**. Thus, even though F_R satisfies **(wC)**, i.e., **(wSP)** for an empty R , it does not for arbitrary R 's. Still, although both F_R and F_M satisfy **(sSP)**, the following result shows that F_M provides a better approximation in terms of property **(SP)**.

Proposition 12

Let P be a program, $V \subseteq \mathcal{A}$, $f \in F_R$, and $f' \in F_M$. Then, for every $R \in \mathcal{C}$ with $\mathcal{A}(R) \subseteq \mathcal{A} \setminus V$,

$$\mathcal{AS}(f(P, V) \cup R) \subseteq \mathcal{AS}(f'(P, V) \cup R).$$

Clearly, F satisfies **(SP)** iff it satisfies **(wSP)** and **(sSP)**. Since no F can in general satisfy **(SP)**, we basically obtain two kinds of relaxations on the conditions of **(SP)**. But we can do even better: following results by Gonçalves et al. (2016b), we know that F satisfies **(SP)** iff it satisfies **(wC)**,

(sC), and (SI). From the results in Props. 1, 8, and 11, we obtain that each of the three discussed classes corresponds to a unique relaxation of the conditions of (SP), implying that our study gives a complete account on which forgetting operators to use when (SP) cannot be satisfied, but only approximated.

Arguably, F_M is also more flexible in situations where we have to forget several atoms for which F_{SP} and F_R do not provide the optimal overall choice.

Example 7

Consider the following program P from which we want to forget about c and p .

$$d \leftarrow c \quad c \leftarrow \text{not not } c \quad a \leftarrow p \quad b \leftarrow \text{not } p \quad p \leftarrow \text{not not } p$$

Clearly, F_{SP} allows us to correctly capture the result of forgetting about c , in the sense that $d \leftarrow \text{not not } d$ is part of the result of forgetting, but, at the same time, will introduce new answer sets in which both a and b are true. On the other hand, F_R will avoid the latter problem, but will simply cancel all rules mentioning d and c . Here, F_M certainly provides the best alternative as it avoids both problems and provides the desired result.

In practice, the choice between the three classes greatly depends on the application at hand. To help making this decision, we now identify, for each of the three classes, a set of conditions in favor of its choice over the other two.

The class F_{SP} should be chosen whenever:

- (SP) should hold for those instances that do not satisfy Ω ;
- Rules that do not mention atoms to be forgotten should be preserved;
- All answer sets should be preserved; and
- We do not mind the appearance of new answer sets.

The class F_R should be chosen whenever:

- Rules that do not mention atoms to be forgotten should be preserved;
- No new answer sets should appear;
- We do not mind that some answer sets may disappear; and
- We do not mind that (SP) does not hold even if Ω does not hold.

The class F_M should be chosen whenever:

- (SP) should hold for those instances that do not satisfy Ω ;
- Answer sets should be preserved precisely (modulo the forgotten atoms); and
- We do not mind to change rules that do not mention atoms to be forgotten.

These conditions directly stem from the properties each of the classes of forgetting operators satisfies, and can be seen as a guideline for a more informed choice between the three alternatives.

6 Complexity

We assume familiarity with standard complexity concepts, such as NP. Given a complexity class \mathcal{C} , a \mathcal{C} oracle decides a given sub-problem from \mathcal{C} in one computation step. The class Σ_k^P contains the problems that can be decided in polynomial time by a non-deterministic Turing machine with unrestricted access to a Σ_{k-1}^P oracle. Π_k^P is the complementary class of Σ_k^P . Thus, $\Sigma_1^P = \text{NP}$, and $\Pi_1^P = \text{coNP}$. We also recall that a language is in complexity class D_i^P iff it is the intersection of a language in Σ_i^P and a language in Π_i^P . Instead of D_1^P we use the more common name D^P . In addition, the following result will be useful due to the established correspondence between HT-and V-HT-models in Prop. 5.

Proposition 13 (Eiter et al. (2007), Theorem 6.12.)

Given a program P , an HT-interpretation $\langle X, Y \rangle$, and $V \subseteq \mathcal{A}$, deciding whether $\langle X, Y \rangle \in \mathcal{HT}_V(P)$ is D^P -complete.

Our first result is in the spirit of model-checking.

Lemma 6.1

Given program P , $V \subseteq \mathcal{A}$, and HT-interpretation $\langle X, Y \rangle$. Deciding whether $\langle X, Y \rangle \in \mathcal{HT}_V(P)_{\parallel V}$ is Σ_2^P -complete. Hardness holds already for disjunctive programs.

Membership follows from guessing an interpretation $Y' \sim_V Y$ and checking $\langle X, Y' \rangle \in \mathcal{HT}_V(P)$ (cf. Proposition 13), while the hardness result can be adapted from the Σ_2^P -hardness of ASP consistency, cf. the paper by Eiter and Gottlob (1995). By means of this, we can determine the complexity of deciding whether a given program is strongly equivalent to the result of forgetting obtained by any $f \in F_R$.

Theorem 4

Given programs P, Q , and $V \subseteq \mathcal{A}$, deciding whether $P \equiv f(Q, V)$ (for $f \in F_R$) is Π_3^P -complete. Hardness holds already for disjunctive programs.

Essentially, for the complementary problem, we guess an HT-interpretation $\langle X, Y \rangle$ and check that either $\langle X, Y \rangle \in \mathcal{HT}(P)$ or $\langle X, Y \rangle \in \mathcal{HT}_V(P)_{\parallel V}$, but not both. The hardness result is then obtained by a reduction from $(3, \forall)$ -QSAT.

The next result provides the complexity of determining whether some X occurs in the intersection of $\mathcal{R}_{\langle P, V \rangle}^Y$ used in the definition of F_{SP} , F_M and Ω .

Lemma 6.2

Given program P , $V \subseteq \mathcal{A}$, and HT-interpretation $\langle X, Y \rangle$ with $Y \subseteq \mathcal{A} \setminus V$, deciding whether $X \in \bigcap \mathcal{R}_{\langle P, V \rangle}^Y$ is in D_2^P .

Basically, we have to perform a Σ_2^P - and a Π_2^P -test. The former decides whether $\mathcal{R}_{\langle P, V \rangle}^Y \neq \emptyset$, while the latter determines that for all $A \subseteq V$, either $\langle Y \cup A, Y \cup A \rangle \notin \mathcal{HT}_V(P)$ or $\langle X, Y \cup A \rangle \in \mathcal{HT}_V(P)$.

This Lemma allows us to obtain an identical result to Thm. 4 for F_{SP} .

Theorem 5

Given programs P, Q , and $V \subseteq \mathcal{A}$, deciding whether $P \equiv f(Q, V)$ (for $f \in F_{SP}$) is Π_3^P -complete. Hardness holds already for disjunctive programs.

The basic proof idea is very similar to the one sketched for Thm. 4, but substituting the test $\langle X, Y \rangle \in \mathcal{HT}_V(P)_{\parallel V}$ with $\langle X, Y \rangle \in \mathcal{HT}(f(P, V))$ for $f \in F_{SP}$.

Since the definition of F_M is based on cases, deciding whether its condition holds, is computationally more expensive than the previous two (in Lemmas 6.1 and 6.2).

Lemma 6.3

Given program P , $V \subseteq \mathcal{A}$, and HT-interpretation $\langle X, Y \rangle$ with $Y \subseteq \mathcal{A} \setminus V$, deciding whether $X \in \bigcup \mathcal{R}_{\langle P, V \rangle}^Y$ if $\mathcal{R}_{\langle P, V \rangle}^Y$ has no least element, and $X \in \bigcap \mathcal{R}_{\langle P, V \rangle}^Y$ otherwise, is in Σ_3^P and in Π_3^P .

Fortunately though, since this test is both in Σ_3^P and in Π_3^P , in the next result, we can basically solve the complementary problem of guessing an HT-interpretation $\langle X, Y \rangle$ and check that either $\langle X, Y \rangle \in \mathcal{HT}(P)$ or $\langle X, Y \rangle \in \mathcal{HT}(f(P, V))$ for $f \in F_M$, but not both, in one step.

Theorem 6

Given programs P , Q , and $V \subseteq \mathcal{A}$, deciding whether $P \equiv f(Q, V)$ (for $f \in F_M$) is Π_3^P -complete. Hardness holds already for disjunctive programs.

Thus, determining whether $P \equiv f(Q, V)$ for f of any of the three considered classes of forgetting operators is always Π_3^P -complete. This shows that the choice which of the three classes of forgetting operators to use in a concrete situation is not influenced by their computational complexity.

Finally, we provide the complexity result for criterion Ω , which on the one hand improves on a flaw for the membership result by Gonçalves et al. (2016c), but also includes the hardness result here.

Theorem 7

Let P be a program over \mathcal{A} and $V \subseteq \mathcal{A}$. Deciding whether $\langle P, V \rangle$ satisfies criterion Ω is Σ_3^P -complete. Hardness holds already for disjunctive programs.

7 Concluding Remarks

We addressed the problem of forgetting in ASP when we *must* forget, even if satisfying the fundamental desirable property (**SP**) is not possible.

We thoroughly investigated three alternatives which, despite stemming from different starting points – one reusing a known class of forgetting operators, one exploring the concept of relativized equivalence, and one trying to get the best of the previous two – turn out to each correspond to the relaxation of one of three properties – (**wC**), (**sC**) and (**SI**) – that together characterize (**SP**). We characterized the three classes by showing which of the usually considered properties each obeys, established links between them, and investigated their computational complexity. The computational complexity turns out to be high, which is not surprising given, for example, the fact that, in classical logic, forgetting can only be expressed as a second-order axiom. Nevertheless, on the one hand, forgetting is an operation not expected to be done as regularly as for example model computation or query answering, while, on the other hand, at least for those classes that satisfy (**SI**), F_{SP} and F_R , we can perform forgetting in a modular way focusing only on the relevant part of the program. Whether this can be extended also to F_M remains an interesting open problem for future research.

We also established relevant novel results concerning a correspondence between V -HT-models and HT-models and a full complexity result for checking whether the criterion (Ω) that indicates whether it is possible to forget while satisfying (**SP**) holds.

It is also noteworthy that there are no other operators or classes of operators mentioned in the literature that satisfy the same set of properties satisfied by any of the three classes discussed here (cf. the paper by Gonçalves et al. (2016b)). The closest approximation is the operator F_{SM} (Wang et al. 2013) which obeys the same set of properties previously found in the literature as F_M , yet, unlike F_M , it does not satisfy either of the inclusions of (**SP**), notably (**sSP**).

Avenues for future research include investigating different forms of forgetting which may be required in practice, such as those that preserve some aggregated meta-level information about the forgotten atoms, or even going beyond maintaining all relationships between non-forgotten atoms which may be required by certain legislation. This may also be of interest for semantics different from ASP, such as for forgetting under the well-founded semantics (Alferes et al. 2013; Knorr and Alferes 2014).

Acknowledgments R. Gonçalves, M. Knorr and J. Leite were partially supported by FCT strategic project NOVA LINC3 (UID/CEC/04516/2013). R. Gonçalves was partially supported by FCT grant SFRH/BPD/100906/2014 and M. Knorr by FCT grant SFRH/BPD/86970/2012. S. Woltran was supported by the Austrian Science Fund (FWF): Y698, P25521.

References

- ALFERES, J. J., KNORR, M., AND WANG, K. 2013. Forgetting under the well-founded semantics. In *Procs. of LPNMR*, P. Cabalar and T. C. Son, Eds. LNCS, vol. 8148. Springer, 36–41.
- 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 (September/October), 43–70.
- BLED3OE, W. W. AND HINES, L. M. 1980. Variable elimination and chaining in a resolution-based prover for inequalities. In *Procs. of CADE*, W. Bibel and R. A. Kowalski, Eds. LNCS, vol. 87. Springer, 70–87.
- DELGRANDE, J. P., SCHAUB, T., TOMPITS, H., AND WOLTRAN, S. 2013. A model-theoretic approach to belief change in answer set programming. *ACM Trans. Comput. Log.* 14, 2, 14.
- DELGRANDE, J. P. AND WANG, K. 2015. A syntax-independent approach to forgetting in disjunctive logic programs. In *Procs. of AAAI*, B. Bonet and S. Koenig, Eds. AAAI Press, 1482–1488.
- 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.
- EITER, T., FINK, M., AND WOLTRAN, S. 2007. Semantical characterizations and complexity of equivalences in answer set programming. *ACM Trans. Comput. Log.* 8, 3, 1–53.
- EITER, T. AND GOTTL0B, G. 1995. On the computational cost of disjunctive logic programming: Propositional case. *Annals of Mathematics and Artificial Intelligence* 15, 3-4, 289–323.
- EITER, T. AND WANG, K. 2008. Semantic forgetting in answer set programming. *Artif. Intell.* 172, 14, 1644–1672.
- EUROPEAN PARLIAMENT. 2016. General data protection regulation. *Official Journal of the European Union* L119/59.
- GONÇALVES, R., KNORR, M., AND LEITE, J. 2016a. Forgetting in ASP: the forgotten properties. In *Procs. of JELIA*, L. Michael and A. C. Kakas, Eds. LNCS, vol. 10021. Springer, 543–550.
- GONÇALVES, R., KNORR, M., AND LEITE, J. 2016b. The ultimate guide to forgetting in answer set programming. In *Procs. of KR*, C. Baral, J. Delgrande, and F. Wolter, Eds. AAAI Press, 135–144.
- GONÇALVES, R., KNORR, M., AND LEITE, J. 2016c. You can’t always forget what you want: on the limits of forgetting in answer set programming. In *Procs. of ECAI*, M. S. Fox and G. A. Kaminka, Eds. IOS Press, 957–965.
- KNORR, M. AND ALFERES, J. J. 2014. Preserving strong equivalence while forgetting. In *Procs. of JELIA*, E. Fermé and J. Leite, Eds. LNCS, vol. 8761. Springer, 412–425.
- KONEV, B., LUDWIG, M., WALTHER, D., AND WOLTER, F. 2012. The logical difference for the lightweight description logic EL. *J. Artif. Intell. Res. (JAIR)* 44, 633–708.
- KONEV, B., LUTZ, C., WALTHER, D., AND WOLTER, F. 2013. Model-theoretic inseparability and modularity of description logic ontologies. *Artif. Intell.* 203, 66–103.
- KONTCHAKOV, R., WOLTER, F., AND ZAKHARYASCHEV, M. 2010. Logic-based ontology comparison and module extraction, with an application to dl-lite. *Artif. Intell.* 174, 15, 1093–1141.
- LANG, J., LIBERATORE, P., AND MARQUIS, P. 2003. Propositional independence: Formula-variable independence and forgetting. *J. Artif. Intell. Res. (JAIR)* 18, 391–443.
- LANG, J. AND MARQUIS, P. 2010. Reasoning under inconsistency: A forgetting-based approach. *Artif. Intell.* 174, 12-13, 799–823.
- LARROSA, J. 2000. Boosting search with variable elimination. In *Procs. of CP*, R. Dechter, Ed. LNCS, vol. 1894. Springer, 291–305.

- LARROSA, J., MORANCHO, E., AND NISO, D. 2005. On the practical use of variable elimination in constraint optimization problems: 'still-life' as a case study. *J. Artif. Intell. Res. (JAIR)* 23, 421–440.
- LEITE, J. 2017. A bird's-eye view of forgetting in answer-set programming. In *Procs. of LPNMR*, M. Balduccini and T. Janhunen, Eds. LNAI, vol. 10377. Springer, 10–22.
- LEWIS, C. I. 1918. *A survey of symbolic logic*. University of California Press. Republished by Dover, 1960.
- LIFSCHITZ, V., PEARCE, D., AND VALVERDE, A. 2001. Strongly equivalent logic programs. *ACM Trans. Comput. Log.* 2, 4, 526–541.
- LIFSCHITZ, V., TANG, L. R., AND TURNER, H. 1999. Nested expressions in logic programs. *Ann. Math. Artif. Intell.* 25, 3-4, 369–389.
- LIN, F. AND REITER, R. 1997. How to progress a database. *Artif. Intell.* 92, 1-2, 131–167.
- LIU, Y. AND WEN, X. 2011. On the progression of knowledge in the situation calculus. In *Procs. of IJCAI*, T. Walsh, Ed. IJCAI/AAAI, 976–982.
- MIDDELDORP, A., OKUI, S., AND IDA, T. 1996. Lazy narrowing: Strong completeness and eager variable elimination. *Theor. Comput. Sci.* 167, 1&2, 95–130.
- MOINARD, Y. 2007. Forgetting literals with varying propositional symbols. *J. Log. Comput.* 17, 5, 955–982.
- RAJARATNAM, D., LEVESQUE, H. J., PAGNUCCO, M., AND THIELSCHER, M. 2014. Forgetting in action. In *Procs. of KR*, C. Baral, G. D. Giacomo, and T. Eiter, Eds. AAAI Press, 498–507.
- 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.
- SLOTA, M. AND LEITE, J. 2012. Robust equivalence models for semantic updates of answer-set programs. In *Procs. of KR*, G. Brewka, T. Eiter, and S. A. McIlraith, Eds. AAAI Press, 158–168.
- SLOTA, M. AND LEITE, J. 2014. The rise and fall of semantic rule updates based on se-models. *TPLP* 14, 6, 869–907.
- WANG, Y., WANG, K., AND ZHANG, M. 2013. Forgetting for answer set programs revisited. In *Procs. of IJCAI*, F. Rossi, Ed. IJCAI/AAAI, 1163–1168.
- WANG, Y., ZHANG, Y., ZHOU, Y., AND ZHANG, M. 2012. Forgetting in logic programs under strong equivalence. In *Procs. of KR*, G. Brewka, T. Eiter, and S. A. McIlraith, Eds. AAAI Press, 643–647.
- WANG, Y., ZHANG, Y., ZHOU, Y., AND ZHANG, M. 2014. Knowledge forgetting in answer set programming. *J. Artif. Intell. Res. (JAIR)* 50, 31–70.
- WANG, Z., WANG, K., TOPOR, R. W., AND PAN, J. Z. 2010. Forgetting for knowledge bases in DL-Lite. *Ann. Math. Artif. Intell.* 58, 1-2, 117–151.
- WEBER, A. 1986. Updating propositional formulas. In *Expert Database Conf.* 487–500.
- WONG, K.-S. 2009. Forgetting in logic programs. Ph.D. thesis, The University of New South Wales.
- ZHANG, Y. AND FOO, N. Y. 2006. Solving logic program conflict through strong and weak forgettings. *Artif. Intell.* 170, 8-9, 739–778.