Updates of Logic Programs by Logic Programs

J. J. Alferes
Dept. Matemática
Univ. Évora and
A.I. Centre
Univ. Nova de Lisboa,
2825 Monte da Caparica
Portugal

J. A. Leite
L. M. Pereira
A.I. Centre
Dept. Informática
Univ. Nova de Lisboa
2825 Monte da Caparica
Portugal

H. Przymusinska
T. Przymusinski
Dept. Computer Science
Univ. of California
Riverside, CA 92521
USA

Abstract

In this paper we investigate updates of knowledge bases represented by logic programs. In order to represent negative information, we use generalized logic programs which allow default negation not only in rule bodies but also in their heads.We start by introducing the notion of an update $P \oplus U$ of a logic program P by another logic program U. Subsequently, we provide a precise semantic characterization of $P \oplus U$, and study some basic properties of program updates. In particular, we show that our update programs generalize the notion of interpretation update.

We then extend this notion to compositional sequences of logic programs updates $P_1 \oplus P_2 \oplus \ldots$, defining a dynamic program update, and thereby introducing the paradigm of *dynamic logic programming*. This paradigm significantly facilitates modularization of logic programming, and thus modularization of non-monotonic reasoning as a whole.

Specifically, suppose that we are given a set of logic program modules, each describing a different state of our knowledge of the world. Different states may represent different time points or different sets of priorities or perhaps even different viewpoints. Consequently, program modules may contain mutually contradictory as well as overlapping information. The role of the dynamic program update is to employ the mutual relationships existing between different modules to precisely determine, at any given module composition stage, the declarative as well as the procedural semantics of the combined program resulting from the modules.

1 Introduction

Most of the work conducted so far in the field of logic programming has focused on representing *static* knowledge, i.e., knowledge that does not evolve with time. This is a serious drawback when dealing with *dynamic knowledge bases* in which not only the *extensional* part (the set of facts) changes dynamically but so does the *intensional* part (the set of rules).

In this paper we investigate updates of knowledge bases represented by logic programs. In order to represent negative information, we use *generalized logic programs* which allow default negation not only in rule bodies but also in their heads. This is needed, in particular, in order to specify that some atoms should became false, i.e., should be deleted. However, our updates are far more expressive than a mere insertion and deletion of facts. They can be specified by means of arbitrary program rules and thus they themselves are logic programs. Consequently, our approach demonstrates how to update one generalized logic program P (the initial program) by another generalized logic program U (the updating program), obtaining as a result a new, updated logic program $P \oplus U$.

Several authors have addressed the issue of updates of logic programs and deductive databases (see e.g. [9, 10, 1]), most of them following the so called *"interpretation update"* approach, originally proposed in [11, 5]. This approach is based on the idea of reducing the problem of finding an update of a knowledge base DB by another knowledge base U to the problem of finding updates of its individual interpretations (models[1]). More precisely, a knowledge base DB' is considered to be the update of a knowledge base DB by U if the set of models of DB'coincides with the set of updated models of DB, i.e. "the set of models of DB'" = "the set of updated models of DB". Thus, according to the interpretation update approach, the problem of finding an update of a *deductive* database DB is reduced to

[1] The notion of a model depends on the type of considered knowledge bases and on their semantics. In this paper we are considering (generalized) logic programs under the stable model semantics.

the problem of finding individual updates of all of its *relational instantiations* (models) M. Unfortunately such an approach suffers, in general, from several important drawbacks[2]:

- In order to obtain the update DB' of a knowledge base DB one has to first compute all the models M of DB (typically, a daunting task) and then individually compute their (possibly multiple) updates M_U by U. An update M_U of a given interpretation M is obtained by changing the status of only those literals in M that are *"forced"* to change by the update U, while keeping all the other literals intact by *inertia* (see e.g. [9, 10]).
- The updated knowledge base DB' is not defined directly but, instead, it is indirectly characterized as a knowledge base whose models coincide with the set of all updated models M_U of DB. In general, there is therefore no natural way of computing[3] DB' because the only straightforward candidate for DB' is the typically intractably large knowledge base DB'' consisting of all clauses that are entailed by all the updated models M_U of DB.
- Most importantly, while the *semantics* of the resulting knowledge base DB' indeed represents the *intended* meaning when just the *extensional* part of the knowledge base DB (the set of facts) is being updated, it leads to strongly *counter-intuitive* results when also the *intensional* part of the database (the set of rules) undergoes change, as the following example shows.

Example 1.1 *Consider the logic program* P :

$$P: \quad sleep \leftarrow not\ tv_on \qquad tv_on \leftarrow$$
$$watch_tv \leftarrow tv_on$$

whose $M = \{tv_on, watch_tv\}$ *is its only stable model. Suppose now that the update* U *states that there is a power failure, and if there is a power failure then the TV is no longer on, as represented by the logic program* U*:*

$$U: \quad not\ tv_on \leftarrow power_failure$$
$$power_failure \leftarrow$$

According to the above mentioned interpretation approach to updating, we would obtain $M_U = \{power_failure, watch_tv\}$ *as the only update of* M *by* U*. This is because* $power_failure$ *needs to be added to the model and its addition forces us to make* tv_on *false. As a result, even though there is a power failure, we are still watching TV. However, by inspecting the initial program and the updating rules, we are likely to conclude that since "watch_tv" was true only because "tv_on" was true, the removal of "tv_on" should make "watch_tv" false by default. Moreover, one would expect "sleep" to become true as well. Consequently, the intended model of the update of* P *by* U *is the model* $M_U' = \{power_failure, sleep\}$.

Suppose now that another update U_2 *follows, described by the logic program:*

$$U_2: \quad not\ power_failure \leftarrow$$

stating that power is back up again. We should now expect the TV to be on again. Since power was restored, i.e. "$power_failure$*" is false, the rule "*$not\ tv_on \leftarrow power_failure$*" of* U *should have no effect and the truth value of "tv_on" should be obtained by inertia from the rule "*$tv_on \leftarrow$*" of the original program* P*.* □

This example illustrates that, when updating knowledge bases, it is not sufficient to just consider the truth values of literals figuring in the heads of its rules because the truth value of their rule bodies may also be affected by the updates of other literals. In other words, it suggests that the *principle of inertia* should be applied not just to the individual literals in an interpretation but rather to entire *rules* of the knowledge base.

The above example also leads us to another important observation, namely, that the notion of an update DB' of one knowledge base DB by another knowledge base U should not just depend on the *semantics* of the knowledge bases DB and U, as it is the case with interpretation updates, but that it should also depend on their *syntax*. This is best illustrated by the following even simpler example:

Example 1.2 *Consider the logic program* P :

$$P: \quad innocent \leftarrow not\ found_guilty$$

[2] In [1] the authors addressed the first two of the drawbacks mentioned below. They showed how to directly construct, given a logic program P, another logic program P' whose partial stable models are exactly the interpretation updates of the partial stable models of P. This eliminates both of these drawbacks (in the case when knowledge bases are logic programs) but it does not eliminate the third, most important drawback.

[3] In fact, in general such a database DB' may not exist at all.

whose only stable model is $M = \{innocent\}$, *because* $found_guilty$ *is false by default. Suppose now that the update* U *states that the person has been found guilty:*

$$U: \quad found_guilty \leftarrow \ .$$

Using the interpretation approach, we would obtain $M_U = \{innocent, found_guilty\}$ *as the only update of* M *by* U *thus leading us to the counter-intuitive conclusion that the person is both innocent and guilty. This is because* $found_guilty$ *must be added to the model* M *and yet its addition does not force us to make innocent false. However, it is intuitively clear that the interpretation* $M'_U = \{found_guilty\}$, *stating that the person is guilty but no longer innocent, should be the only model of the updated program. Observe, however, that the program* P *is* semantically equivalent *to the following program* P' :

$$P': \quad innocent \leftarrow$$

because the programs P *and* P' *have exactly the same set of stable models, namely the model* M. *Nevertheless, while the model* $M_U = \{innocent, found_guilty\}$ *is not the intended model of the update of* P *by* U *it is in fact the only reasonable model of the update of* P' *by* U. □

In this paper we investigate the problem of updating knowledge bases represented by generalized logic programs and we propose a new solution to this problem that attempts to eliminate the drawbacks of the previously proposed approaches. Specifically, given one generalized logic program P (the so called initial program) and another logic program U (the updating program) we define a new generalized logic program $P \oplus U$ called the *update* of P by U. The definition of the updated program $P \oplus U$ does not require any computation of the models of either P or U and is in fact obtained by means of a simple, *linear-time* transformation of the programs P and U. As a result, the update transformation can be accomplished very efficiently and its *implementation* is quite straightforward[4].

Due to the fact that we apply the inertia principle not just to atoms but to entire program rules, the semantics of our updated program $P \oplus U$ avoids the drawbacks of interpretation updates and moreover it seems to properly represent the intended semantics. As mentioned above, the updated program $P \oplus U$ does not just depend on the *semantics* of the programs P and U, as it was the case with interpretation updates, but it also depends on their *syntax*. In order to make the meaning of the updated program clear and easily verifiable, we provide a *complete characterization* of the semantics of updated programs $P \oplus U$.

Nevertheless, while our notion of program update significantly differs from the notion of interpretation update, it coincides with the latter (as originally introduced in [9] under the name of *revision program* and later reformulated in the language of logic programs in [10]) when the initial program P is purely *extensional*, i.e., when the initial program is just a set of facts. Our definition also allows significant flexibility and can be easily modified to handle updates which incorporate *contradiction removal* or specify different inertia rules. Consequently, our approach can be viewed as introducing a general dynamic logic programming *framework* for updating programs which can be suitably modified to make it fit different application domains and requirements.

Finally, we extend the notion of program updates to sequences of programs, defining the so called *dynamic program updates*. The idea of dynamic updates is very simple and quite fundamental. Suppose that we are given a set of program modules P_s, indexed by different states of the world s. Each program P_s contains some knowledge that is supposed to be true at the state s. Different states may represent different time periods or different sets of priorities or perhaps even different viewpoints. Consequently, the individual program modules may contain mutually contradictory as well as overlapping information. The role of the dynamic program update $\bigoplus \{P_s : s \in S\}$ is to use the mutual relationships existing between different states (as specified by the order relation) to precisely determine, at any given state s, the *declarative* as well as the *procedural* semantics of the combined program, composed of all modules.

Consequently, the notion of a dynamic program update supports the important paradigm of *dynamic logic programming*. Given individual and largely *independent* program modules P_s describing our knowledge at different states of the world (for example, the knowledge acquired at different times), the dynamic program update $\bigoplus \{P_s : s \in S\}$ specifies the exact meaning of the union of these programs. Dynamic programming significantly facilitates *modularization* of logic programming and, thus, modularization of non-monotonic reasoning as a whole. Whereas traditional logic programming has concerned itself mostly with representing static knowledge, we show how to use logic programs to represent dynamically changing knowledge.

Our results extend and improve upon the approach initially proposed in [7], where the authors first argued that the principle of inertia should be applied to the rules of the initial program rather than to the individual literals in

[4] The implementation is available from: `http://www-ssdi.di.fct.unl.pt/~jja/updates/`.

an interpretation. However, the specific update transformation presented in [7] suffered from some drawbacks and was not sufficiently general.

We begin in Section 2 by defining the language of *generalized logic programs,* which allow default negation in rule heads. We describe stable model semantics of such programs as a special case of the approach proposed earlier in [8]. In Section 3 we define the program update $P \oplus U$ of the initial program P by the updating program U. In Section 4 we provide a complete characterization of the semantics of program updates $P \oplus U$ and in Section 5 we study their basic properties. In Section 6 we introduce the notion of dynamic program updates. We close the paper with concluding remarks and notes on future research.

2 Generalized Logic Programs and their Stable Models

In order to represent *negative* information in logic programs and in their updates, we need more general logic programs that allow default negation $not\ A$ not only in premises of their clauses but also in their heads.[5]. We call such programs *generalized logic programs.* In this section we introduce generalized logic programs and extend the stable model semantics of normal logic programs [3] to this broader class of programs[6].

The class of generalized logic programs can be viewed as a special case of a yet broader class of programs introduced earlier in [8]. While our definition is different and seems to be simpler than the one used in [8], when restricted to the language that we are considering, the two definitions can be shown to be equivalent[7].

It will be convenient to *syntactically* represent generalized logic programs as *propositional Horn theories.* In particular, we will represent default negation $not\ A$ as a standard propositional variable (atom). Suppose that $\mathcal{K}$ is an arbitrary set of propositional variables whose names do not begin with a "*not*". By the propositional language $\mathcal{L}_\mathcal{K}$ *generated* by the set $\mathcal{K}$ we mean the language $\mathcal{L}$ whose set of propositional variables consists of:

$$\{A : A \in \mathcal{K}\} \cup \{not\ \mathcal{A} : \mathcal{A} \in \mathcal{K}\}.$$

Atoms $A \in \mathcal{K}$, are called *objective atoms* while the atoms $not\ A$ are called *default atoms.* From the definition it follows that the two sets are disjoint.

By a *generalized logic program* P in the language $\mathcal{L}_\mathcal{K}$ we mean a finite or infinite set of propositional Horn clauses of the form:

$$L \leftarrow L_1, \ldots, L_n$$

where L and L_i are atoms from $\mathcal{L}_\mathcal{K}$. If all the atoms L appearing in heads of clauses of P are objective atoms, then we say that the logic program P is *normal.* Consequently, from a syntactic standpoint, a logic program is simply viewed as a propositional Horn theory. However, its *semantics* significantly differs from the semantics of classical propositional theories and is determined by the class of stable models defined below.

By a (2-valued) *interpretation* M of $\mathcal{L}_\mathcal{K}$ we mean any set of atoms from $\mathcal{L}_\mathcal{K}$ that satisfies the condition that for any A in $\mathcal{K}$, precisely one of the atoms A or $not\ A$ belongs to M. Given an interpretation M we define:

$$\begin{aligned} M^+ &= \{A \in \mathcal{K} : \mathcal{A} \in \mathcal{M}\} \\ M^- &= \{not\ A : not\ A \in M\} = \\ &= \{\ not\ A : A \notin M\}. \end{aligned}$$

Definition 2.1 (Stable models of generalized logic programs) *We say that a (2-valued) interpretation M of $\mathcal{L}_\mathcal{K}$ is a stable model of a generalized logic program P if M is the least model of the Horn theory $P \cup M^-$:*

$$M = Least(P \cup M^-),$$

or, equivalently, if $M = \{L : L$ is an atom and $P \cup M^- \vdash L\}$. □

Example 2.1 *Consider the program:*

$$\begin{array}{ccc} a \leftarrow not\ b & c \leftarrow b & e \leftarrow not\ d \\ not\ d \leftarrow not\ c, a & d \leftarrow not\ e & \end{array}$$

[5] For further motivation and intuitive reading of logic programs with default negations in the heads see [8].

[6] In a forthcoming paper we extend our results to 3-valued (partial) models of logic programs, and, in particular, to well-founded models.

[7] Note that the class of generalized logic programs differs from the class of programs with the so called *"classical"* negation [4] which allow the use of *strong* rather than default negation in their heads.

and let $\mathcal{K} = \{\dashv, \lfloor, \rfloor, \lceil, \rceil\}$. *This program has precisely one stable model* $M = \{a, e, not\, b, not\, c, not\, d\}$. *To see that* M *is stable we simply observe that:*

$$M = Least(P \cup \{not\, b, not\, c, not\, d\}).$$

The interpretation $N = \{not\, a, not\, e, b, c, d\}$ *is not a stable model because:*

$$N \neq Least(P \cup \{not\, e, not\, a\}). \ \square$$

Following an established tradition, from now on we will be omitting the default (negative) atoms when describing interpretations and models. Thus the above model M will be simply listed as $M = \{a, e\}$. The following Proposition easily follows from the definition of stable models.

Proposition 2.1 *An interpretation* M *of* $\mathcal{L}_\mathcal{K}$ *is a stable model of a generalized logic program* P *if and only if*

$$M^+ = \{A : A \in \mathcal{K} \text{ and } \frac{\mathcal{P}}{\mathcal{M}} \vdash \mathcal{A}\}$$

and

$$M^- \supseteq \{not\, A : A \in \mathcal{K} \text{ and } \frac{\mathcal{P}}{\mathcal{M}} \vdash not\mathcal{A}\},$$

where $\frac{P}{M}$ *denotes the Gelfond-Lifschitz transform [3] of* P *w.r.t.* M. $\square$

Clearly, the second condition in the above Proposition is always vacuously satisfied for normal programs and therefore we immediately obtain:

Proposition 2.2 *The class of stable models of generalized logic programs extends the class of stable models of normal programs [3].* $\square$

3 Program Updates

Suppose that $\mathcal{K}$ is an arbitrary set of propositional variables, and P and U are two generalized logic programs in the language $\mathcal{L} = \mathcal{L}_\mathcal{K}$. By $\widehat{\mathcal{K}}$ we denote the following superset of $\mathcal{K}$:

$$\widehat{\mathcal{K}} = \mathcal{K} \cup \{\mathcal{A}^-,\ \mathcal{A}_\mathcal{P},\ \mathcal{A}^-_\mathcal{P},\ \mathcal{A}_\mathcal{U},\ \mathcal{A}^-_\mathcal{U} : \mathcal{A} \in \mathcal{K}\}.$$

This definition assumes that the original set $\mathcal{K}$ of propositional variables does not contain any of the newly added symbols of the form $A^-, A_P, A^-_P, A_U, A^-_U$ so that they are all disjoint sets of symbols. If $\mathcal{K}$ contains any such symbols then they have to be *renamed* before the extension of $\mathcal{K}$ takes place. We denote by $\widehat{\mathcal{L}} = \mathcal{L}_{\widehat{\mathcal{K}}}$ the extension of the language $\mathcal{L} = \mathcal{L}_\mathcal{K}$ generated by $\widehat{\mathcal{K}}$.

Definition 3.1 (Program Updates) *Let* P *and* U *be generalized programs in the language* $\mathcal{L}$. *We call* P *the original program and* U *the updating program. By the update of* P *by* U *we mean the generalized logic program* $P \oplus U$, *which consists of the following clauses in the extended language* $\widehat{\mathcal{L}}$*:*

(RP) Rewritten original program clauses:

$$A_P \leftarrow B_1, \ldots, B_m, C^-_1, \ldots, C^-_n \quad (1)$$
$$A^-_P \leftarrow B_1, \ldots, B_m, C^-_1, \ldots, C^-_n \quad (2)$$

for any clause:

$$A \leftarrow B_1,\ \ldots,\ B_m,\ not\, C_1,\ \ldots,\ not\, C_n$$

and

$$not\, A \leftarrow B_1,\ \ldots,\ B_m, not\, C_1,\ \ldots,\ not\, C_n$$

respectively, in the original program P. *The rewritten clauses are obtained from the original ones by replacing atoms* A *(respectively, the atoms* $not\, A$*) occurring in their heads by the atoms* A_P *(respectively,* A^-_P*) and by replacing negative premises* $not\, C$ *by* C^-.

(RU) Rewritten updating program clauses:

$$A_U \leftarrow B_1, \ldots, B_m, C_1^-, \ldots, C_n^- \quad (3)$$
$$A_U^- \leftarrow B_1, \ldots, B_m, C_1^-, \ldots, C_n^- \quad (4)$$

for any clause:

$$A \leftarrow B_1, \ \ldots, \ B_m, \ not\, C_1, \ \ldots, \ not\, C_n$$

and, respectively,

$$not\, A \leftarrow B_1, \ \ldots, \ B_m, not\, C_1, \ \ldots, \ not\, C_n$$

in the updating program U. The rewritten clauses are obtained from the original ones by replacing atoms A (respectively, the atoms $not\, A$) occurring in their heads by the atoms A_U (respectively, A_U^-) and by replacing negative premises $not\, C$ by C^-.

(UR) Update rules:

$$A \leftarrow A_U \qquad A^- \leftarrow A_U^- \quad (5)$$

for all objective atoms $A \in \mathcal{K}$. The update rules state that an atom A must be true (respectively, false) in $P \oplus U$ if it is true (respectively, false) in the updating program U.

(IR) Inheritance rules:

$$A \leftarrow A_P, not\, A_U^- \qquad A^- \leftarrow A_P^-, not\, A_U \quad (6)$$

for all objective atoms $A \in \mathcal{K}$. The inheritance rules say that an atom A (respectively, A^-) in $P \oplus U$ is inherited (by inertia) from the original program P provided it is not rejected (i.e., forced to be false) by the updating program U. More precisely, an atom A is true (respectively, false) in $P \oplus U$ if it is true (respectively, false) in the original program P, provided it is not made false (respectively, true) by the updating program U.

(DR) Default rules:

$$A^- \leftarrow not\, A_P, not\, A_U \qquad not\, A \leftarrow A^- \quad (7)$$

for all objective atoms $A \in \mathcal{K}$. The first default rule states that an atom A in $P \oplus U$ is false if it is neither true in the original program P nor in the updating program U. The second says that if an atom is false then it can be assumed to be false by default. It ensures that A and A^- cannot both be true. □

It is easy to show that any model N of $P \oplus U$ is *coherent*, i.e., A is true (respectively, false) in N iff A^- is false (respectively, true) in N, for any $A \in \mathcal{K}$. In other words, every stable model of $P \oplus U$ satisfies the constraint $not\, A \equiv A^-$. Consequently, A^- can be simply regarded as an internal (meta-level) representation of the default negation $not\, A$ of A.

Example 3.1 *Consider the programs P and U from Example 1.1:*

$$P: \quad sleep \leftarrow not\, tv_on \qquad tv_on \leftarrow$$
$$watch_tv \leftarrow tv_on$$

$$U: \quad not\ tv_on \leftarrow power_failure$$
$$power_failure \leftarrow$$

The update of the program P by the program U is the logic program $P \oplus U = (RP) \cup (RU) \cup (UR) \cup (IR) \cup (DR)$, where:

$$RP: \quad sleep_P \leftarrow tv_on^- \qquad tv_on_P \leftarrow$$
$$watch_tv_P \leftarrow tv_on$$

$$RU: \quad tv_on_U^- \leftarrow power_failure$$
$$power_failure_U \leftarrow$$

It is easy to verify that $M = \{power_failure, sleep\}$ is the only stable model (modulo irrelevant literals) of $P \oplus U$. □

4 Semantic Characterization of Program Updates

In this section we provide a complete semantic characterization of update programs $P \oplus U$ by describing their stable models. This characterization shows precisely how the semantics of the update program $P \oplus U$ depends on the syntax and semantics of the programs P and U.

Let P and U be *fixed* generalized logic programs in the language $\mathcal{L}$. Since the update program $P \oplus U$ is defined in the extended language $\widehat{\mathcal{L}}$, we begin by showing how interpretations of the language $\mathcal{L}$ can be extended to interpretations of the extended language $\widehat{\mathcal{L}}$.

Definition 4.1 (Extended Interpretation) *For any interpretation M of $\mathcal{L}$ we denote by $\widehat{M}$ its extension to an interpretation of the extended language $\widehat{\mathcal{L}}$ defined, for any atom $A \in \mathcal{K}$, by the following rules:*

$$\begin{array}{lll} A^- \in \widehat{M} & \textit{iff} & \textit{not } A \in M \\ A_P \in \widehat{M} & \textit{iff} & \exists\ A \leftarrow \textit{Body} \in P \ \textit{ and } \ M \models \textit{Body} \\ A_P^- \in \widehat{M} & \textit{iff} & \exists\ \textit{not } A \leftarrow \textit{Body} \in P \\ & & \textit{and } \ M \models \textit{Body} \end{array}$$

$$\begin{array}{lll} A_U \in \widehat{M} & \textit{iff} & \exists A \leftarrow \textit{Body} \in U \textit{ and } \ M \models \textit{Body} \\ A_U^- \in \widehat{M} & \textit{iff} & \exists\ \textit{not } A \leftarrow \textit{Body} \in U \\ & & \textit{and } \ M \models \textit{Body}. \end{array}$$ □

We will also need the following definition:

Definition 4.2 *For any model M of the program U in the language $\mathcal{L}$ define:*

$$\begin{array}{l} \textit{Defaults}[M] = \\ \{\textit{not } \ A : M \models \neg \textit{Body}, \forall (A \ \leftarrow \ \textit{Body}) \in P \ \cup \ U\}; \end{array}$$

$$\begin{array}{l} \textit{Rejected}[M] = \\ \{A \leftarrow \ \textit{Body} \in P : \exists\ (\textit{not } \ A \ \leftarrow \ \textit{Body}' \in U) \\ \quad \textit{and } \ M \models \textit{Body}'\} \\ \cup \\ \{\textit{not } \ A \leftarrow \ \textit{Body} \in P : \exists\ (A \ \leftarrow \ \textit{Body}' \in U) \\ \quad \textit{and } \ M \models \textit{Body}'\}; \end{array}$$

$$\textit{Residue}[M] = P \ \cup \ U - \textit{Rejected}[M].$$ □

The set $\textit{Defaults}[M]$ contains default negations $\textit{not } A$ of all *unsupported* atoms A, i.e., atoms that have the property that the body of every clause from $P \cup U$ with the head A is false in M. Consequently, negation $\textit{not } A$ of these unsupported atoms A can be assumed by default. The set $\textit{Rejected}[M] \subseteq P$ represents the set of clauses of the original program P that are *rejected* (or contradicted) by the update program U and its model M. The residue $\textit{Residue}[M]$ consists of all clauses in the union $P \cup U$ of programs P and U that were *not* rejected by the update program U. Note that all the three sets depend on the model M as well as on the *syntax* of the programs P and U.

Now we are able to describe the semantics of the update program $P \oplus U$ by providing a complete characterization of its stable models.

Theorem 4.1 (Characterization of stable models of update programs) *An interpretation N of the language $\widehat{\mathcal{L}} = \mathcal{L}_{\widehat{\mathcal{K}}}$ is a stable model of the update $P \oplus U$ if and only if N is the extension $N = \widehat{M}$ of a model M of U that satisfies the condition:*

$$M = \textit{Least}(P \ \cup \ U - \textit{Rejected}[M] \ \cup \ \textit{Defaults}[M]),$$

or $M = \textit{Least}(\textit{Residue}[M] \ \cup \ \textit{Defaults}[M])$, equivalently. □

Example 4.1 *Consider again the programs P and U from Example 1.1. Let $M = \{\textit{power_failure}, \textit{sleep}\}$. We obtain:*

$$\begin{array}{l} \textit{Defaults}[M] = \{\textit{not watch_tv}, \} \\ \textit{Rejected}[M] = \{\textit{tv_on} \leftarrow\} \\ \textit{Residue}[M] = \left\{ \begin{array}{l} \textit{sleep} \leftarrow \textit{not tv_on} \\ \textit{watch_tv} \leftarrow \textit{tv_on} \\ \textit{not tv_on} \leftarrow \textit{power_failure} \\ \textit{power_failure} \leftarrow \end{array} \right\} \end{array}$$

and thus it is easy to see that

$$M = Least(Residue[M] \cup Defaults[M]).$$

Consequently, $\widehat{M}$ *is a stable model of the update program* $P \oplus U$. □

5 Properties of Program Updates

In this section we study the basic properties of program updates. Since $Defaults[M] \subseteq M^-$, we conclude that the condition $M = Least(Residue[M] \cup Defaults[M])$ clearly implies $M = Least(Residue[M] \cup M^-)$ and thus we immediately obtain:

Proposition 5.1 *If* N *is a stable model of* $P \oplus U$ *then its restriction* $M = N|\mathcal{L}$ *to the language* $\mathcal{L}$ *is a stable model of* $Residue[M]$. □

However, the condition $M = Least(Residue[M] \cup Defaults[M])$ says much more than just that M is a stable model of $Residue[M]$. It says that M is completely *determined* by the set $Defaults[M]$, i.e., by the set of negations of unsupported atoms that can be assumed false by default.

Clearly, if M is a stable model of $P \cup U$ then $Rejected[M] = \emptyset$ and $Defaults[M] = M^-$, which implies:

Proposition 5.2 *If* M *is a stable model of the union* $P \cup U$ *of programs* P *and* U *then its extension* $N = \widehat{M}$ *is a stable model of the update program* $P \oplus U$. *Thus, the semantics of the update program* $P \oplus U$ *is always weaker than or equal to the semantics of the union* $P \cup U$ *of programs* P *and* U. □

In general, the converse of the above result does not hold. In particular, the union $P \cup U$ may be a contradictory program with no stable models.

Example 5.1 *Consider again the programs* P *and* U *from Example 1.1. It is easy to see that* $P \cup U$ *is contradictory.* □

If either P or U is empty and M is a stable model of $P \cup U$ then $Rejected[M] = \emptyset$ and therefore M is also a stable model of $P \oplus U$.

Proposition 5.3 *If either* P *or* U *is empty then* M *is a stable model of* $P \cup U$ *iff* $N = \widehat{M}$ *is a stable model of* $P \oplus U$. *Thus, in this case, the semantics of the update program* $P \oplus U$ *coincides with the semantics of the union* $P \cup U$. □

Proposition 5.4 *If both* P *and* U *are normal programs (or if both have only clauses with default atoms* $not\ A$ *in their heads) then* M *is a stable model of* $P \cup U$ *iff* $N = \widehat{M}$ *is a stable model of* $P \oplus U$. *Thus, in this case the semantics of the update program* $P \oplus U$ *also coincides with the semantics of the union* $P \cup U$ *of programs* P *and* U. □

5.1 Program Updates Generalize Interpretation Updates

In this section we show that *interpretation updates*, originally introduced under the name *"revision programs"* by Marek and Truszczynski [9], and subsequently given a simpler characterization by Przymusinski and Turner [10], constitute a special case of program updates. Here, we identify the *"revision rules"*:

$$\begin{array}{rcl} in(A) & \leftarrow & in(B), out(C) \\ out(A) & \leftarrow & in(B), out(C) \end{array}$$

used in [9], with the following generalized logic program clauses:

$$\begin{array}{rcl} A & \leftarrow & B, not\ C \\ not\ A & \leftarrow & B, not\ C. \end{array}$$

Theorem 5.1 (Program updates generalize interpretation updates) *Let* I *be any interpretation and* U *any updating program in the language* $\mathcal{L}$. *Denote by* P_I *the generalized logic program in* $\mathcal{L}$ *defined by*

$$P_I = \{A \leftarrow : A \in I\} \cup \{not\ A \leftarrow : not\ A \in I\}.$$

Then $\widehat{J}$ *is a stable model of the program update* $P_I \oplus U$ *of the program* P_I *by the program* U *iff* J *is an interpretation update of* I *by* U *(in the sense of [9]).* □

This theorem shows that when the initial program P is purely *extensional*, i.e., contains only positive or negative *facts*, then the interpretation update of P by U is semantically equivalent to the updated program $P \oplus U$. As shown by the Examples 1.1 and 1.2, when P contains deductive rules then the two notions become significantly different.

6 Dynamic Program Updates

In this section we introduce the notion of *dynamic program update* $\bigoplus\{ P_s : s \in S\}$ over an ordered set $\mathcal{P} = \{ P_s : s \in S\}$ of logic programs which provides an important generalization of the notion of single program updates $P \oplus U$ introduced in Section 3.

The idea of dynamic updates, inspired by [6], is simple and quite fundamental. Suppose that we are given a set of program modules P_s, indexed by different states of the world s. Each program P_s contains some knowledge that is supposed to be true at the state s. Different states may represent different time periods or different sets of priorities or perhaps even different viewpoints. Consequently, the individual program modules may contain mutually contradictory as well as overlapping information. The role of the dynamic program update $\bigoplus\{P_s : s \in S\}$ is to use the mutual relationships existing between different states (and specified in the form of the ordering relation) to precisely determine, at any given state s, the *declarative* as well as the *procedural* semantics of the combined program, composed of all modules.

Consequently, the notion of a dynamic program update supports the important paradigm of *dynamic logic programming*. Given individual and largely *independent* program modules P_s describing our knowledge at different states of the world (for example, the knowledge acquired at different times), the dynamic program update $\bigoplus\{P_s : s \in S\}$ specifies the exact meaning of the union of these programs. Dynamic programming significantly facilitates modularization of logic programming and, thus, modularization of non-monotonic reasoning as a whole.

Suppose that $\mathcal{P} = \{\mathcal{P}_\int : \int \in \mathcal{S}\}$ is a finite or infinite sequence of generalized logic programs in the language $\mathcal{L} = \mathcal{L}_\mathcal{K}$, indexed by the set $S = \{1, 2, \ldots, n, \ldots\}$. We will call elements s of the set $S \cup \{0\}$ *states* and we will refer to 0 as the *initial state*. If S has the *largest* element then we will denote it by *max*.

Remark 6.1 *Instead of a linear sequence of states $S \cup \{0\}$ one could as well consider any finite or infinite ordered set with the smallest element s_0 and with the property that every state s other than s_0 has an immediate predecessor $s-1$ and that $s_0 = s - n$, for some finite n. In particular, one may use a finite or infinite tree with the root s_0 and the property that every node (state) has only a finite number of ancestors.* □

By $\overline{\mathcal{K}}$ we denote the following superset of the set $\mathcal{K}$ of propositional variables:

$$\overline{\mathcal{K}} = \mathcal{K} \cup \{ \quad A^-, A_s,\ A_s^-, A_{P_s},\ A_{P_s}^-, reject(A_s), \\ reject(A_s^-) : A \in \mathcal{K},\ \int \in \mathcal{S} \cup \{\imath\}\}.$$

As before, this definition assumes that the original set $\mathcal{K}$ of propositional variables does not contain any of the newly added symbols of the form $A^-, A_s,\ A_s^-, A_{P_s},\ A_{P_s}^-$) so that they are all disjoint sets of symbols. If the original language $\mathcal{K}$ contains any such symbols then they have to be *renamed* before the extension of $\mathcal{K}$ takes place. We denote by $\overline{\mathcal{L}} = \mathcal{L}_{\overline{\mathcal{K}}}$ the extension of the language $\mathcal{L} = \mathcal{L}_\mathcal{K}$ generated by $\overline{\mathcal{K}}$.

Definition 6.1 (Dynamic Program Update) *By the dynamic program update over the sequence of updating programs $\mathcal{P} = \{\mathcal{P}_\int : \int \in \mathcal{S}\}$ we mean the logic program $\biguplus\mathcal{P}$, which consists of the following clauses in the extended language $\overline{\mathcal{L}}$:*

(RP) Rewritten program clauses:

$$A_{P_s} \leftarrow B_1, \ldots, B_m, C_1^-, \ldots, C_n^- \tag{8}$$
$$A_{P_s}^- \leftarrow B_1, \ldots, B_m, C_1^-, \ldots, C_n^- \tag{9}$$

for any clause:

$$A \leftarrow B_1,\ \ldots,\ B_m,\ not\, C_1,\ \ldots,\ not\, C_n$$

respectively, for any clause:

$$not\, A \leftarrow B_1,\ \ldots,\ B_m, not\, C_1,\ \ldots,\ not\, C_n$$

in the program P_s, where $s \in S$. The rewritten clauses are simply obtained from the original ones by replacing atoms A (respectively, the atoms $not\, A$) occurring in their heads by the atoms A_{P_s} (respectively, $A_{P_s}^-$) and by replacing negative premises $not\, C$ by C^-.

(UR) Update rules:

$$A_s \leftarrow A_{P_s}; \qquad A_s^- \leftarrow A_{P_s}^- \tag{10}$$

for all objective atoms $A \in \mathcal{K}$ and for all $s \in S$. The update rules state that an atom A must be true (respectively, false) in the state $s \in S$ if it is true (respectively, false) in the updating program P_s.

(IR) Inheritance rules:

$$A_s \leftarrow A_{s-1}, not\ A^-_{P_s};$$

$$A^-_s \leftarrow A^-_{s-1}, not\ A_{P_s}$$

for all objective atoms $A \in \mathcal{K}$ and for all $s \in S$. The inheritance rules say that an atom A is true (respectively, false) in the state $s \in S$ if it is true (respectively, false) in the previous state $s-1$ and it is not rejected, *i.e., forced to be false (respectively, true), by the updating program P_s.*

(DR) Default rules (describing the initial state):

$$A^-_0, \tag{11}$$

for all objective atoms $A \in \mathcal{K}$. Default rules describe the initial state 0 by making all objective atoms initially false. □

Observe that the dynamic program update $\biguplus\mathcal{P}$ is a normal logic program, i.e., it does not contain default negation in heads of its clauses. Moreover, only the inheritance rules contain default negation in their bodies. Also note that the program $\biguplus\mathcal{P}$ does not contain the atoms A or A^-, where $A \in \mathcal{K}$, in heads of its clauses. These atoms appear only in the bodies of rewritten program clauses. The notion of the dynamic program update $\bigoplus_s\mathcal{P}$ at a given state $s \in S$ changes that.

Definition 6.2 (Dynamic Program Update at a Given State) *Given a fixed state $s \in S$, by the dynamic program update at the state s, denoted by $\bigoplus_s\mathcal{P}$, we mean the dynamic program update $\biguplus\mathcal{P}$ augmented with the following:*

Current State Rules CS(s):

$$A \leftarrow A_s \qquad A^- \leftarrow A^-_s \qquad not\ A \leftarrow A^-_s \tag{12}$$

for all objective atoms $A \in \mathcal{K}$. Current state rules specify the current state s in which the updated program is being evaluated and determine the values of the atoms A, A^- and not A. In particular, if the set S has the largest element max *then we simply write $\bigoplus\mathcal{P}$ instead of $\bigoplus_{max}\mathcal{P}$.* □

Mark that whereas for any state s $\biguplus\mathcal{P}$ is not required to be coherent, $\bigoplus_s\mathcal{P}$ must be so.

The notion of a dynamic program update generalizes the previously introduced notion of an update $P \oplus U$ of two programs P and U.

Theorem 6.1 *Let P_1 and P_2 be arbitrary generalized logic programs and let $S = \{1, 2\}$. The dynamic program update $\bigoplus\{P_1, P_2\} = \bigoplus_2\{P_1, P_2\}$ at the state $max = 2$ is semantically equivalent to the program update $P_1 \oplus P_2$ defined in Section 3.* □

6.1 Examples

Example 6.1 *Let $\mathcal{P} = \{P_1, P_2, P_3\}$, where P_1, P_2 and P_3 are as follows:*

P_1 : $sleep \leftarrow not\ tv_on$
$watch_tv \leftarrow tv_on$
$tv_on \leftarrow$

P_2 : $not\ tv_on \leftarrow power_failure$
$power_failure \leftarrow$

P_3 : $not\ power_failure \leftarrow$

The dynamic program update over $\mathcal{P}$ is the logic program $\biguplus\mathcal{P} = (RP_1) \cup (RP_2) \cup (RP_3) \cup (UR) \cup (IR) \cup (DR)$, where

$$
\begin{array}{ll}
RP_1: & sleep_{P_1} \leftarrow tv_on^- \\
& watch_tv_{P_1} \leftarrow tv_on \\
& tv_on_{P_1} \leftarrow \\
\\
RP_2: & tv_on^-_{P_2} \leftarrow power_failure \\
& power_failure_{P_2} \leftarrow \\
\\
RP_3: & power_failure^-_{P_3} \leftarrow
\end{array}
$$

and the dynamic program update at the state s is $\bigoplus_s \mathcal{P} = \biguplus\mathcal{P} \cup \mathcal{CS}(\int)$. Consequently, as intended, $\bigoplus_1\mathcal{P}$ has a single stable model $M_1 = \{tv_on, watch_tv\}$; $\bigoplus_2\mathcal{P}$ has a single stable model $M_2 = \{sleep, power_failure\}$ and $\bigoplus\mathcal{P} = \bigoplus_3\mathcal{P}$ has a single stable model $M_3 = \{tv_on, watch_tv\}$ (all models modulo irrelevant literals). Moreover. $\bigoplus_2\mathcal{P}$ is semantically equivalent to $P_1 \oplus P_2$. □

As mentioned in the Introduction, in dynamic logic programming, logic program modules describe states of our knowledge of the world, where different states may represent different time points or different sets of priorities or even different viewpoints. It is not our purpose in this paper to discuss in detail how to apply dynamic logic programming to any of these application domains[8]. However, since all of the examples presented so far relate different program modules with changing time, below we illustrate how to use dynamic logic programming to represent the well known problem in the domain of taxonomies by using priorities among rules.

Example 6.2 *Consider the well-known problem of flying birds. In this example we have several rules with different priorities. First, the animals-do-not-fly rule, which has the lowest priority; then the birds-fly rule with a higher priority; the penguins-do-not-fly rule with an even higher priority; and, finally, with the highest priority, all the rules describing the actual taxonomy (penguins are birds, birds are animals, etc). This can be coded quite naturally in dynamic logic programming:*

$$
\begin{array}{ll}
P_1: & not\ fly(X) \leftarrow animal(X) \\
P_2: & fly(X) \leftarrow bird(X) \\
P_3: & not\ fly(X) \leftarrow penguin(X) \\
P_4: & animal(X) \leftarrow bird(X) \\
& bird(X) \leftarrow penguin(X) \\
& animal(pluto) \\
& bird(duffy) \\
& penguin(tweety)
\end{array}
$$

The reader can check that, as intended, the dynamic logic program at state 4, i.e. $\bigoplus_4\{P_1, P_2, P_3, P_4\}$, has a single stable model where $fly(duffy)$ is true, and both $fly(pluto)$ and $fly(tweety)$ are false. □

7 Conclusions and Future Work

We defined a program transformation that takes two generalized logic programs P and U, and produces the updated logic program $P \oplus U$ resulting from the update of program P by U. We provided a complete characterization of the semantics of program updates $P \oplus U$ and we established their basic properties. Our approach generalizes the so called *revision programs* introduced in [9]. Namely, in the special case when the initial program is just a set of facts, our program update coincides with the justified revision of [9]. In the general case, when the initial program also contains rules, our program updates characterize precisely which of these rules remain valid by inertia, and which are rejected.

With the introduction of dynamic program updates, we have extended program updates to ordered sets of logic programs (or modules). When this order is interpreted as a time order, dynamic program updates describe the evolution of a logic program which undergoes a sequence of modifications. This opens up the possibility of incremental design and evolution of logic programs, leading to the paradigm of *dynamic logic programming*. We believe that dynamic programming significantly facilitates *modularization* of logic programming, and, thus, modularization of non-monotonic reasoning as a whole.

[8] In fact, this is the subject of our ongoing research. In particular, the application of dynamic logic programming to the domain of actions is the subject of a forthcoming paper.

A specific application of dynamic logic programming that we intend to explore, is the evolution and maintenance of software specifications. By using logic programming as a specification language, dynamic programming provides the means of representing the evolution of software specifications.

However, ordered sets of program modules need not necessarily be seen as just a temporal evolution of a logic program. Different modules can also represent different sets of priorities, or viewpoints of different agents. In the case of priorities, a dynamic program update specifies the exact meaning of the "union" of the modules, subject to the given priorities. We intend to further study the relationship between dynamic logic programming and other preference-based approaches to knowledge representation.

Although not explored in here, a dynamic program update can be queried not only about the current state but also about other states. If modules are seen as viewpoints of different agents, the truth of some A_s in $\bigoplus \mathcal{P}$ can be read as: A is true according to agent s in a situation where the knowledge of the $\bigoplus \mathcal{P}$ is "visible" to agent s.

We are in the process of generalizing our approach and results to the *3-valued* case, which will enable us to update programs under the well-founded semantics. We have already developed a working implementation for the 3-valued case with top-down querying.

Our approach to program updates has grown out of our research on representing non-monotonic knowledge by means of logic programs. We envisage enriching it in the near future with other dynamic programming features, such as abduction and contradiction removal. Among other applications that we intend to study are productions systems modelling, reasoning about concurrent actions and active and temporal databases.

Acknowledgements

This work was partially supported by PRAXIS XXI project MENTAL, by JNICT project ACROPOLE, by the National Science Foundation grant # IRI931-3061, and a NATO scholarship while the L. M. Pereira was on sabbatical leave at the Department of Computer Science, University of California, Riverside. The work of J. A. Leite was supported by PRAXIS Scholarship no. BD/13514/97.

References

[1] J. J. Alferes, L. M. Pereira. *Update-programs can update programs.* In J. Dix, L. M. Pereira and T. Przymusinski, editors, Selected papers from the ICLP'96 Workshop NMELP'96, vol. 1216 of LNAI, pages 110-131. Springer-Verlag, 1997.

[2] J. J. Alferes, L. M. Pereira and T. Przymusinski. *Strong and Explicit Negation in Non-Monotonic Reasoning and Logic Programming.* In J. J. Alferes, L. M. Pereira and E. Orlowska, editors, JELIA '96, volume 1126 of LNAI, pages 143-163. Springer-Verlag, 1996.

[3] M. Gelfond and V. Lifschitz. *The stable model semantics for logic programming.* In R. Kowalski and K. A. Bowen. editors. 5th International Logic Programming Conference, pages 1070-1080. MIT Press, 1988.

[4] M. Gelfond and V. Lifschitz. *Logic Programs with classical negation.* In Warren and Szeredi, editors, 7th International Logic Programming Conference, pages 579-597. MIT Press, 1990.

[5] H. Katsuno and A. Mendelzon. *On the difference between updating a knowledge base and revising it.* In James Allen, Richard Fikes and Erik Sandewall, editors, Principles of Knowledge Representation and Reasoning: Proceedings of the Second International Conference (KR91), pages 230-237, Morgan Kaufmann 1991.

[6] João A. Leite. *Logic Program Updates.* M.Sc. Dissertation, Universidade Nova de Lisboa, 1997.

[7] J. A. Leite and L. M. Pereira. *Generalizing updates: from models to programs.* In LPKR'97: ILPS'97 Workshop on Logic Programming and Knowledge Representation, Port Jefferson, NY, USA, October 13-16, 1997.

[8] V. Lifschitz and T. Woo. *Answer sets in general non-monotonic reasoning (preliminary report).* In B. Nebel, C. Rich and W. Swartout, editors, Principles of Knowledge Representation and Reasoning, Proceedings of the Third International Conference (KR92), pages 603-614. Morgan-Kaufmann, 1992

[9] V. Marek and M. Truszczynski. *Revision specifications by means of programs.* In C. MacNish, D. Pearce and L. M. Pereira, editors, JELIA '94, volume 838 of LNAI, pages 122-136. Springer-Verlag, 1994.

[10] T. Przymusinski and H. Turner. *Update by means of inference rules.* In V. Marek, A. Nerode, and M. Truszczynski, editors, LPNMR'95, volume 928 of LNAI, pages 156-174. Springer-Verlag, 1995.

[11] M. Winslett. *Reasoning about action using a possible models approach.* In Proceeding of AAAI'88, pages 89-93. 1988.