

A Comparative Study of Logic Programs with Preference

Torsten Schaub* and Kewen Wang

Institut für Informatik, Universität Potsdam
Postfach 60 15 53, D-14415 Potsdam, Germany
{torsten,kewen}@cs.uni-potsdam.de

Abstract

We are interested in semantical underpinnings for existing approaches to preference handling in extended logic programming (within the framework of answer set programming). As a starting point, we explore three different approaches that have been recently proposed in the literature. Because these approaches use rather different formal means, we furnish a series of uniform characterizations that allow us to gain insights into the relationships among these approaches. To be more precise, we provide different characterizations in terms of (i) fixpoints, (ii) order preservation, and (iii) translations into standard logic programs. While the two former provide semantics for logic programming with preference information, the latter furnishes implementation techniques for these approaches.

1 Introduction

Numerous approaches to logic programming with preference information have been proposed in the literature. So far, however, there is no systematic account on their structural differences, finally leading to solid semantical underpinnings. We address this shortcoming by a comparative study of a distinguished class of approaches to preference handling. This class consists of *selective* approaches remaining within the complexity class of extended logic programming (under answer sets semantics). These approaches are selective insofar as they use preferences to distinguish certain “models” of the original program.

We explore three different approaches that have been recently proposed in the literature, namely the ones in [Brewka and Eiter, 1999; Delgrande *et al.*, 2000; Wang *et al.*, 2000]. Our investigation adopts characterization techniques found in the same literature in order to shed light on the relationships among these approaches. This provides us with different characterizations in terms of (i) fixpoints, (ii) order preservation, and (iii) translations into standard logic programs. While the two former provide semantics for logic programming with preference information, the latter furnishes implementation techniques for these approaches. From another

* Affiliated with the School of Computing Science at Simon Fraser University, Burnaby, Canada.

perspective, one can view (iii) as an axiomatization of the underlying strategy within the object language, while (i) may be regarded as a meta-level description of the corresponding construction process. One may view (ii) as the most semantical characterization because it tells us which “models” of the original program are selected by the respective preference handling strategy.

We limit (also in view of (iii)) our investigation to approaches to preference handling that remain within NP. This excludes approach like the ones in [Rintanen, 1995; Zhang and Foo, 1997] that step outside the complexity class of the underlying reasoning method. This applies also to the approach in [Sakama and Inoue, 1996], where preferences on literals are investigated. While the approach of [Gelfond and Son, 1997] remains within NP, it advocates strategies that are non-selective. Approaches that can be addressed within this framework include [Baader and Hollunder, 1993; Brewka, 1994] that were originally proposed for default logic.

2 Definitions and notation

We assume a basic familiarity with logic programming under *answer set semantics* [Gelfond and Lifschitz, 1991]. An *extended logic program* is a finite set of rules of the form

$$L_0 \leftarrow L_1, \dots, L_m, \text{not } L_{m+1}, \dots, \text{not } L_n, \quad (1)$$

where $n \geq m \geq 0$, and each L_i ($0 \leq i \leq n$) is a *literal*, ie. either an atom A or its negation $\neg A$. The set of all literals is denoted by *Lit*. Given a rule r as in (1), we let $\text{head}(r)$ denote the *head*, L_0 , of r and $\text{body}(r)$ the *body*, $\{L_1, \dots, L_m, \text{not } L_{m+1}, \dots, \text{not } L_n\}$, of r . Further, let $\text{body}^+(r) = \{L_1, \dots, L_m\}$ and $\text{body}^-(r) = \{L_{m+1}, \dots, L_n\}$. A program is called *basic* if $\text{body}^-(r) = \emptyset$ for all its rules.

We define the *reduct* of a rule r as $r^+ = \text{head}(r) \leftarrow \text{body}^+(r)$. The *reduct*, Π^X , of a program Π *relative to* a set X of literals is defined by

$$\Pi^X = \{r^+ \mid r \in \Pi \text{ and } \text{body}^-(r) \cap X = \emptyset\}.$$

A set of literals X is *closed under* a basic program Π iff for any $r \in \Pi$, $\text{head}(r) \in X$ whenever $\text{body}^+(r) \subseteq X$. We say that X is *logically closed* iff it is either consistent (ie. it does not contain both a literal A and its negation $\neg A$) or equals *Lit*. The smallest set of literals which is both logically closed and closed under a basic program Π is denoted by $Cn(\Pi)$.

Finally, a set X of literals is an *answer set* of a program Π iff $Cn(\Pi^X) = X$. In what follows, we deal with *consistent* answer sets only.

The set Γ_{Π}^X of all *generating rules* of an answer set X from Π is given by

$$\Gamma_{\Pi}^X = \{r \in \Pi \mid \text{body}^+(r) \subseteq X \text{ and } \text{body}^-(r) \cap X = \emptyset\}.$$

As van Gelder in [1993], we define $C_{\Pi}(X) = Cn(\Pi^X)$. Note that the operator C_{Π} is anti-monotonic, which implies that the operator $A_{\Pi}(X) = C_{\Pi}(C_{\Pi}(X))$ is monotonic. A fixpoint of A_{Π} is called an *alternating fixpoint* for Π . Different semantics are captured by distinguishing different groups of fixpoints of A_{Π} .

A (statically) *ordered logic program*¹ is a pair $(\Pi, <)$, where Π is an extended logic program and $< \subseteq \Pi \times \Pi$ is an irreflexive and transitive relation. Given, $r_1, r_2 \in \Pi$, the relation $r_1 < r_2$ expresses that r_2 has higher priority than r_1 .²

3 Preferred alternating fixpoints

The notion of answer sets (without preference) is based on a reduction of extended logic programs to basic programs (without default negation). Such a reduction is inapplicable when addressing conflicts by means of preference information since all conflicts between rules are simultaneously resolved when turning Π into Π^X . Rather conflict resolution must be characterized among the original rules in order to account for blockage between rules. That is, once the negative body $\text{body}^-(r)$ is eliminated there is no way to detect whether $\text{head}(r') \in \text{body}^-(r)$ holds in case of $r < r'$.

Such an approach is pursued in [Wang et al., 2000] for characterizing “preferred” answer sets. Following earlier approaches based on default logic [Baader and Hollunder, 1993; Brewka, 1994], this approach is based on the concept of *activeness*: Let $X, Y \subseteq \text{Lit}$ be two sets of literals in an ordered logic program $(\Pi, <)$. A rule r in Π is *active* wrt the pair (X, Y) , if $\text{body}^+(r) \subseteq X$ and $\text{body}^-(r) \cap Y = \emptyset$.

Definition 1 (Wang et al., 2000) Let $(\Pi, <)$ be an ordered logic program and let X be a set of literals. We define

$$\begin{aligned} X_0 &= \emptyset \quad \text{and for } i \geq 0 \\ X_{i+1} &= X_i \cup \{ \text{head}(r) \mid \left. \begin{array}{l} \text{I. } r \in \Pi \text{ is active wrt } (X_i, X) \text{ and} \\ \text{II. there is no rule } r' \in \Pi \text{ with } r < r' \\ \text{such that} \\ \quad (a) \text{ } r' \text{ is active wrt } (X, X_i) \text{ and} \\ \quad (b) \text{ } \text{head}(r') \notin X_i \end{array} \right\} \end{aligned}$$

Then, $C_{(\Pi, <)}(X) = \bigcup_{i \geq 0} X_i$ if $\bigcup_{i \geq 0} X_i$ is consistent.

Otherwise, $C_{(\Pi, <)}(X) = \text{Lit}$.

The idea is to apply a rule r only if the question of application has been settled for all higher-ranked rules r' . That is, if either its prerequisites will never be derivable, viz. $\text{body}^+(r') \not\subseteq X$, or r' is defeated by what has been derived so far, viz.

¹Also called *prioritized* logic program by some authors, as eg. in [Brewka and Eiter, 1999].

²Some authors, among them [Brewka and Eiter, 1999], attribute relation $<$ the inverse meaning.

$\text{body}^-(r) \cap X_i \neq \emptyset$, or r' or another rule with the same head have already applied, viz. $\text{head}(r') \in X_i$.

As its original C_{Π} , the operator $C_{(\Pi, <)}$ is anti-monotonic. Accordingly, we may define for any set $X \subseteq \text{Lit}$, the *alternating transformation* of $(\Pi, <)$ as $\mathcal{A}_{(\Pi, <)}(X) = C_{(\Pi, <)}(C_{(\Pi, <)}(X))$. A fixpoint of $\mathcal{A}_{(\Pi, <)}$ is called an *alternating fixpoint* of $(\Pi, <)$. Note that $\mathcal{A}_{(\Pi, <)}$ is monotonic.

Now, in analogy to Van Gelder [1993], a semantical framework for ordered logic programs in terms of sets of alternating fixpoints can be defined. Three different types of semantics are investigated in [Wang et al., 2000]: (i) Preferred³ answer sets, viz. alternating fixpoints being also fixpoints of $C_{(\Pi, <)}$. (ii) Preferred regular extensions, viz. maximal normal⁴ alternating fixpoints of $(\Pi, <)$. (iii) Preferred well-founded model, viz. the least alternating fixpoint of $(\Pi, <)$.

We put the prefix ‘w-’ whenever a distinction to other approaches is necessary.

For illustration, consider the following ordered logic program $(\Pi_2, <)$ due to [Baader and Hollunder, 1993]:

$$\begin{array}{lll} r_1 : & \neg f & \leftarrow p, \text{not } f & r_2 < r_1 \\ r_2 : & w & \leftarrow b, \text{not } \neg w \\ r_3 : & f & \leftarrow w, \text{not } \neg f \\ r_4 : & b & \leftarrow p \\ r_5 : & p & \leftarrow \end{array} \quad (2)$$

Observe that Π_2 admits two answer sets: $X = \{p, b, \neg f, w\}$ and $X' = \{p, b, f, w\}$. As argued in [Baader and Hollunder, 1993], X is the unique w-preferred answer set. To see this, observe that

$$\begin{array}{ll} X_0 = \emptyset & X'_0 = \emptyset \\ X_1 = \{p\} & X'_1 = \{p\} \\ X_2 = \{p, b, \neg f\} & X'_2 = \{p, b\} \\ X_3 = \{p, b, \neg f, w\} & X'_3 = X'_2 \neq X' \\ X_4 = X_3 = X & \end{array}$$

Note that w cannot be included into X'_3 because r_1 is active wrt (X', X'_2) and r_1 is preferred to r_2 .

4 Compiling order preservation

A translation of ordered logic programs $(\Pi, <)$ to standard ones Π' is developed in [Delgrande et al., 2000]. The specific strategy used there ensures that the resulting program Π' admits only those answer sets of the original program Π that are *order preserving*:

Definition 2 Let $(\Pi, <)$ be a statically ordered program and let X be an answer set of Π .

Then, X is called *<-preserving*, if there exists an enumeration $\langle r_i \rangle_{i \in I}$ of Γ_{Π}^X such that for every $i, j \in I$ we have that:

0. $\text{body}^+(r_i) \subseteq \{\text{head}(r_j) \mid j < i\}$; and
1. if $r_i < r_j$, then $j < i$; and
2. if $r_i < r'$ and $r' \in \Pi \setminus \Gamma_{\Pi}^X$, then
 - (a) $\text{body}^+(r') \not\subseteq X$ or
 - (b) $\text{body}^-(r') \cap \{\text{head}(r_j) \mid j < i\} \neq \emptyset$.

³Originally called *prioritized*.

⁴An alternating fixpoint X is normal if $X \subseteq C_{(\Pi, <)}(X)$.

Condition 0 makes the property of *groundedness*⁵ explicit. Although any standard answer set is generated by a grounded sequence of rules, we will see in the sequel that this property is weakened when preferences are at issue. Condition 1 stipulates that $\langle r_i \rangle_{i \in I}$ is *compatible* with $<$, a property invariant to all of the considered approaches. Lastly, Condition 2 is comparable with Condition II in Definition 1; it guarantees that rules can never be blocked by lower-ranked rules.

As above, $X = \{p, b, \neg f, w\}$ is the only $<$ -preserving answer set of Π_2 ; it can be generated by the grounded sequences $\langle r_5, r_4, r_1, r_2 \rangle$ and $\langle r_5, r_1, r_4, r_2 \rangle$ both of which satisfy conditions 1 and 2. The only grounded sequence generating $X' = \{p, b, f, w\}$, namely $\langle r_5, r_4, r_2, r_3 \rangle$, violates 2b.

The corresponding translation integrates ordering information into the logic program via a special-purpose predicate symbol $<$. This allows also for treating ordering information in a dynamic fashion. A logic program over a propositional language $\mathcal{L}$ is said to be *dynamically ordered* iff $\mathcal{L}$ contains the following pairwise disjoint categories: (i) a set N of terms serving as *names* for rules; (ii) a set At of (propositional) atoms of a program; and (iii) a set $At_{<}$ of *preference atoms* $s < t$, where $s, t \in N$ are names. For each such program Π , we assume furthermore a bijective function $n(\cdot)$ assigning to each rule $r \in \Pi$ a name $n(r) \in N$. To simplify notation, we usually write n_r instead of $n(r)$ (and we sometimes abbreviate n_{r_i} by n_i).

An atom $n_r < n_{r'}$ in $At_{<}$ amounts to asserting that $r < r'$ holds. A statically ordered program $(\Pi, <)$ can thus be captured by programs containing preference atoms only among their facts; it is then expressed by the program $\Pi \cup \{(n_r < n_{r'}) \leftarrow r < r'\}$.

Given $r < r'$, one wants to ensure that r' is considered before r , in the sense that, for a given answer set X , rule r' is known to be applied or defeated *ahead of* r (cf. Condition II or 2 above, respectively). This is done by translating rules so that the order of rule application can be explicitly controlled. For this purpose, one needs to be able to detect when a rule has been applied or when a rule is defeated. For a rule r , there are two cases for it not to be applied: it may be that some literal in $body^+(r)$ does not appear in the answer set, or it may be that a literal in $body^-(r)$ is in the answer set. For detecting non-applicability (i.e., blockage), for each rule r in the given program Π , a new, special-purpose atom $bl(n_r)$ is introduced. Similarly, a special-purpose atom $ap(n_r)$ is introduced to detect the case where a rule has been applied. For controlling application of rule r the atom $ok(n_r)$ is introduced. Informally, one concludes that it is ok to apply a rule just if it is ok with respect to every $<$ -greater rule; for such a $<$ -greater rule r' , this will be the case just when r' is known to be blocked or applied.

More formally, given a dynamically ordered program Π over $\mathcal{L}$, let $\mathcal{L}^+$ be the language obtained from $\mathcal{L}$ by adding, for each $r, r' \in \Pi$, new pairwise distinct propositional atoms $ap(n_r)$, $bl(n_r)$, $ok(n_r)$, and $ok'(n_r, n_{r'})$. Then, the translation $\mathcal{T}$ maps an ordered program Π over $\mathcal{L}$ into a standard program $\mathcal{T}(\Pi)$ over $\mathcal{L}^+$ in the following way.

⁵This term is borrowed from the literature on default logic.

Definition 3 (Delgrande et al., 2000) Let $\Pi = \{r_1, \dots, r_k\}$ be a dynamically ordered logic program over $\mathcal{L}$.

Then, the logic program $\mathcal{T}(\Pi)$ over $\mathcal{L}^+$ is defined as $\mathcal{T}(\Pi) = \bigcup_{r \in \Pi} \tau(r)$, where $\tau(r)$ consists of the following rules, for $L^+ \in body^+(r)$, $L^- \in body^-(r)$, and $r', r'' \in \Pi$:

$$\begin{aligned} a_1(r) : \quad & head(r) \leftarrow ap(n_r) \\ a_2(r) : \quad & ap(n_r) \leftarrow ok(n_r), body(r) \\ b_1(r, L^+) : \quad & bl(n_r) \leftarrow ok(n_r), not L^+ \\ b_2(r, L^-) : \quad & bl(n_r) \leftarrow ok(n_r), L^- \\ c_1(r) : \quad & ok(n_r) \leftarrow ok'(n_r, n_{r_1}), \dots, ok'(n_r, n_{r_k}) \\ c_2(r, r') : \quad & ok'(n_r, n_{r'}) \leftarrow not (n_r < n_{r'}) \\ c_3(r, r') : \quad & ok'(n_r, n_{r'}) \leftarrow (n_r < n_{r'}), ap(n_{r'}) \\ c_4(r, r') : \quad & ok'(n_r, n_{r'}) \leftarrow (n_r < n_{r'}), bl(n_{r'}) \\ t(r, r', r'') : \quad & n_r < n_{r''} \leftarrow n_r < n_{r'}, n_{r'} < n_{r''} \\ as(r, r') : \quad & \neg(n_r < n_{r'}) \leftarrow n_r < n_{r'} \end{aligned}$$

We write $\mathcal{T}(\Pi, <)$ rather than $\mathcal{T}(\Pi')$, whenever Π' is the dynamically ordered program capturing $(\Pi, <)$.

The first four rules of $\tau(r)$ express applicability and blocking conditions of the original rules. The second group of rules encodes the strategy for handling preferences. The first of these rules, $c_1(r)$, “quantifies” over the rules in Π . This is necessary when dealing with dynamic preferences since preferences may vary depending on the corresponding answer set. The three rules $c_2(r, r')$, $c_3(r, r')$, and $c_4(r, r')$ specify the pairwise dependency of rules in view of the given preference ordering: For any pair of rules r, r' with $n_r < n_{r'}$, we derive $ok'(n_r, n_{r'})$ whenever $n_r < n_{r'}$ fails to hold, or whenever either $ap(n_{r'})$ or $bl(n_{r'})$ is true. This allows us to derive $ok(n_r)$, indicating that r may potentially be applied whenever we have for all r' with $n_r < n_{r'}$ that r' has been applied or cannot be applied. It is important to note that this is only one of many strategies for dealing with preferences: different strategies are obtainable by changing the specification of $ok(\cdot)$ and $ok'(\cdot, \cdot)$, as we will see below.

As shown in [Delgrande et al., 2000], a set of literals X is a $<$ -preserving answer set of Π iff $X = Y \cap \mathcal{L}$ for some answer set Y of $\mathcal{T}(\Pi, <)$. In the sequel, we refer to such answer sets as being *D-preferred*.

5 Synthesis

The last two sections have exposed three rather different ways of characterizing preferred answer sets. Despite their different characterizations, however, it turns out that the two approaches prefer similar answer sets.

5.1 Characterizing D-preference

We start by providing a fixpoint definition for *D-preference*. For this purpose, we assume a bijective mapping $rule(\cdot)$ from rule heads to rules, that is, $rule(head(r)) = r$; accordingly, $rule(\{head(r) \mid r \in R\}) = R$. Such mappings can be defined in a bijective way by distinguishing different occurrences of literals.

Definition 4 Let $(\Pi, <)$ be a statically ordered logic pro-

gram and let X be a set of literals. We define

$$\begin{aligned} X_0 &= \emptyset \quad \text{and for } i \geq 0 \\ X_{i+1} &= X_i \cup \{ \text{head}(r) \mid \\ &\quad \left. \begin{array}{l} \text{I. } r \in \Pi \text{ is active wrt } (X_i, X) \text{ and} \\ \text{II. there is no rule } r' \in \Pi \text{ with } r < r' \\ \text{such that} \\ \quad (a) r' \text{ is active wrt } (X, X_i) \text{ and} \\ \quad (b) r' \notin \text{rule}(X_i) \end{array} \right\} \end{aligned}$$

Then, $\mathcal{C}_{(\Pi, <)}^D(X) = \bigcup_{i \geq 0} X_i$ if $\bigcup_{i \geq 0} X_i$ is consistent.

Otherwise, $\mathcal{C}_{(\Pi, <)}^D(X) = \text{Lit}$.

The difference between this definition and Definition 1 manifests itself in IIb. While D-preference requires that a higher-ranked rule has effectively applied, w-preference contents itself with the presence of the head of the rule, no matter whether this was supplied by the rule itself.

This difference is nicely illustrated by program $(\Pi_3, <)$:

$$\begin{array}{ll} r_1 : a \leftarrow \text{not } b & r_2 < r_1 \\ r_2 : b \leftarrow & \\ r_3 : a \leftarrow & \end{array} \quad (3)$$

While the only answer set $\{a, b\}$ is w-preferred set, there is no D-preferred answer set. This is the same with program $(\Pi'_3, <)$ obtained by replacing r_1 with $r'_1 : a \leftarrow b$.

We have the following result providing three alternative characterizations of D-preferred answer sets.

Theorem 1 Let $(\Pi, <)$ be a statically ordered logic program over $\mathcal{L}$ and let X be a consistent set of literals.

Then, the following propositions are equivalent.

1. $\mathcal{C}_{(\Pi, <)}^D(X) = X$;
2. $X = Y \cap \mathcal{L}$ for some answer set Y of $\mathcal{T}(\Pi, <)$;
3. X is a $<$ -preserving answer set of Π .

While the last result dealt with effective answer sets, the next one shows that applying operator $\mathcal{C}_{(\Pi, <)}^D$ is equivalent to the application of van Gelder's operator $\mathcal{C}_{\Pi'}$ to the translated program $\mathcal{T}(\Pi, <)$.

Theorem 2 Let $(\Pi, <)$ be a statically ordered logic program over $\mathcal{L}$ and let X be a consistent set of literals over $\mathcal{L}$.

Then, we have that $\mathcal{C}_{(\Pi, <)}^D(X) = \mathcal{C}_{\mathcal{T}(\Pi, <)}(Y) \cap \mathcal{L}$ for some set of literals Y over $\mathcal{L}^+$ such that $X = Y \cap \mathcal{L}$.

This result is important because it allows us to use the translation $\mathcal{T}(\Pi, <)$ for implementing further semantics by appeal to the alternating fixpoint idea.

5.2 Characterizing w-preference

We start by showing how w-preference can be characterized in terms of order preservation.

Definition 5 Let $(\Pi, <)$ be a statically ordered program and let X be an answer set of Π .

Then, X is called $<^w$ -preserving, if there exists an enumeration $\langle r_i \rangle_{i \in I}$ of Γ_{Π}^X such that for every $i, j \in I$ we have that:

0. (a) $\text{body}^+(r_i) \subseteq \{\text{head}(r_j) \mid j < i\}$ or

(b) $\text{head}(r_i) \in \{\text{head}(r_j) \mid j < i\}$; and

1. if $r_i < r_j$, then $j < i$; and

2. if $r_i < r'$ and $r' \in \Pi \setminus \Gamma_{\Pi}^X$, then

(a) $\text{body}^+(r') \not\subseteq X$ or

(b) $\text{body}^-(r') \cap \{\text{head}(r_j) \mid j < i\} \neq \emptyset$ or

(c) $\text{head}(r') \in \{\text{head}(r_j) \mid j < i\}$.

The primary difference of this concept of order preservation to the original one is clearly the weaker notion of groundedness. This involves the rules in Γ_{Π}^X (via Condition 0b) as well as those in $\Pi \setminus \Gamma_{\Pi}^X$ (via Condition 2c). The rest of the definition is the same as in Definition 2. For instance, answer set $\{a, b\}$ of Π_3 is generated by the $<^w$ -preserving rule sequence $\langle r_3, r_2 \rangle$. Note that r_1 satisfies 2c but neither 2a nor 2b. For a complement, in $(\Pi'_3, <)$, r'_1 is dealt with via Condition 0b.

Interestingly, this weaker notion of groundedness can be easily integrated into the translation given in the last section.

Definition 6 Given the same prerequisites as in Definition 3.

Then, the logic program $\mathcal{T}^w(\Pi)$ over $\mathcal{L}^+$ is defined as $\mathcal{T}^w(\Pi) = \bigcup_{r \in \Pi} \tau(r) \cup \{c_5(r, r') \mid r, r' \in \Pi\}$, where

$$c_5(r, r') : \text{ok}'(n_r, n_{r'}) \leftarrow (n_r < n_{r'}), \text{head}(r')$$

The purpose of $c_5(r, r')$ is to eliminate rules from the preference handling process once their head has been derived.

We have the following result, showing in particular, how w-preference is implementable via off-the-shelf logic programming systems.

Theorem 3 Let $(\Pi, <)$ be a statically ordered logic program over $\mathcal{L}$ and let X be a consistent set of literals. Then, the following propositions are equivalent.

1. $\mathcal{C}_{(\Pi, <)}(X) = X$;
2. $X = Y \cap \mathcal{L}$ for some answer set Y of $\mathcal{T}^w(\Pi, <)$;
3. X is a $<^w$ -preserving answer set of Π .

In analogy to what we have shown above, we have the following stronger result, opening the avenue for implementing more semantics based on w-preference:

Theorem 4 Let $(\Pi, <)$ be a statically ordered logic program over $\mathcal{L}$ and let X be a consistent set of literals over $\mathcal{L}$.

Then, we have that $\mathcal{C}_{(\Pi, <)}(X) = \mathcal{C}_{\mathcal{T}^w(\Pi, <)}(Y) \cap \mathcal{L}$ for some set of literals Y over $\mathcal{L}^+$ such that $X = Y \cap \mathcal{L}$.

6 Brewka and Eiter's concept of preference

Another approach to preference was proposed by Brewka and Eiter in [1999]. For brevity, we omit technical details and simply say that an answer set is B-preferred; the reader is referred to [Brewka and Eiter, 1999; 2000] for details.

This approach differs in two significant ways from the two approaches given above. First, the construction of answer sets is separated from verifying whether they respect the given preferences. Interestingly, this verification is done on the basis of the prerequisite-free program obtained from the original one by "evaluating" $\text{body}^+(r)$ for each rule r wrt the separately constructed (standard) answer set. Second, rules that putatively lead to counter-intuitive results are removed from

the inference process. This is made explicit in [Brewka and Eiter, 2000], where the following filtering transformation is defined:⁶

$$Z_X(\Pi) = \Pi \setminus \{r \in \Pi \mid \text{head}(r) \in X, \text{body}^-(r) \cap X \neq \emptyset\} \quad (4)$$

Then, by definition, an answer set of Π is B-preferred iff it is a B-preferred answer set of $Z_X(\Pi)$.

The distinguishing example of this approach is given by program $(\Pi_5, <)$:

$$\begin{array}{llll} r_1 : & b & \leftarrow & a, \text{not } \neg b & r_3 < r_2 < r_1 \\ r_2 : & \neg b & \leftarrow & \text{not } b \\ r_3 : & a & \leftarrow & \text{not } \neg a \end{array} \quad (5)$$

Program Π_5 has two standard answer sets, $\{a, b\}$ and $\{a, \neg b\}$. While the former is B-preferred, neither of them is w- or D-preferred (see below). Also, we note that both answer sets of program $(\Pi_2, <)$ are B-preferred, while only $\{p, b, \neg f, w\}$ is w- and D-preferred.

In order to shed some light on these differences, we start by providing a fixpoint characterization of B-preference:

Definition 7 Let $(\Pi, <)$ be an ordered logic program and let X be a set of literals. We define

$$\begin{aligned} X_0 &= \emptyset & \text{and for } i \geq 0 \\ X_{i+1} &= X_i \cup \{ \text{head}(r) \mid \left. \begin{array}{l} \text{I. } r \in \Pi \text{ is active wrt } (X, X) \text{ and} \\ \text{II. there is no rule } r' \in \Pi \text{ with } r < r' \\ \text{such that} \\ \quad (a) \ r' \text{ is active wrt } (X, X_i) \text{ and} \\ \quad (b) \ \text{head}(r') \notin X_i \end{array} \right\} \end{aligned}$$

Then, $C_{(\Pi, <)}^B(X) = \bigcup_{i \geq 0} X_i$ if $\bigcup_{i \geq 0} X_i$ is consistent.

Otherwise, $C_{(\Pi, <)}(X) = \text{Lit}$.

The difference between this definition⁷ and its predecessors manifests itself in Condition I, where activeness is tested wrt (X, X) instead of (X_i, X) as in Definition 1 and 4. In fact, in Example (5) it is the (unprovability of the) prerequisite a of the highest-ranked rule r_1 that makes the construction of w- or D-preferred answer sets break down (cf. Definition 1 and 4). This is avoided with B-preference because once answer set $\{a, b\}$ is provided, its preference-compatibility is tested wrt the program obtained by replacing r_1 with $b \leftarrow \text{not } \neg b$.

B-preference can be captured by means of the following notion of order preservation:

Definition 8 Let $(\Pi, <)$ be a statically ordered program and let X be an answer set of Π .

Then, X is called $<^B$ -preserving, if there exists an enumeration $\langle r_i \rangle_{i \in I}$ of Γ_{Π}^X such that, for every $i, j \in I$, we have that:

1. if $r_i < r_j$, then $j < i$; and

⁶While this is integrated into [Brewka and Eiter, 1999, Def. 4.4], it is made explicit in [Brewka and Eiter, 2000, Def. 6].

⁷We have refrained from integrating (4) in order to keep the fixpoint operator comparable to those given in the previous sections. This is taken care of in the second proposition of Theorem 5.

2. if $r_i < r'$ and $r' \in \Pi \setminus \Gamma_{\Pi}^X$, then

- (a) $\text{body}^+(r') \not\subseteq X$ or
- (b) $\text{body}^-(r') \cap \{\text{head}(r_j) \mid j < i\} \neq \emptyset$ or
- (c) $\text{head}(r') \in X$.

This definition differs in two ways from its predecessors. First, it drops any requirement on groundedness, expressed by Condition 0 above. This corresponds to using (X, X) instead of (X_i, X) in Definition 7. Hence, groundedness is fully disconnected from order preservation. In fact, observe that the B-preferred answer set $\{a, b\}$ of $(\Pi_5, <)$ is associated with the $<^B$ -preserving sequence $\langle r_1, r_2 \rangle$, while the standard answer set itself is generated by the grounded sequence $\langle r_2, r_1 \rangle$.

Second, Condition 2c is more relaxed than in Definition 5. That is, any rule r' whose head is in X (as opposed to X_i) is taken as “applied”. Apart from this, Condition 2c also integrates the filter-conditions from (4).⁸ For illustration, consider Example (3) extended by $r_3 < r_2$:

$$\begin{array}{llll} r_1 : & a & \leftarrow & \text{not } b & r_3 < r_2 < r_1 \\ r_2 : & b & \leftarrow & \\ r_3 : & a & \leftarrow & \end{array} \quad (6)$$

While this program has no D- or w-preferred answer set, it has a B-preferred one: $\{a, b\}$ generated by $\langle r_2, r_3 \rangle$. The critical rule r_1 is handled by 2c. As a net result, Condition 2 is weaker than its counterpart in Definition 5.

We have the following results.

Theorem 5 Let $(\Pi, <)$ be a statically ordered logic program over $\mathcal{L}$ and let X be a consistent answer set of Π .

Then, the following propositions are equivalent.

1. X is B-preferred;
2. $C_{(Z_X(\Pi), <)}^B(X) = X$;
3. $X = Y \cap \mathcal{L}$ for some answer set Y of $\mathcal{T}^B(\Pi, <)$ (where $\mathcal{T}^B$ is defined in [Delgrande et al., 2000]);
4. X is a $<^B$ -preserving answer set of Π .

Unlike theorems 1 and 3, the last result stipulates that X must be an answer set of Π . This requirement can only be dropped in case 3, while all other cases rely on this property.

7 Relationships

Up to now, we have tried to clarify the structural differences between the respective approaches. This has led to homogeneous characterizations that allow us to compare the examined approaches in a uniform way. As a result, we obtain insights into the relationships among these approaches.

First of all, we observe that all three approaches treat the blockage of (higher-ranked) rules in the same way. That is, a rule r' is found to be blocked if either its prerequisites in $\text{body}^+(r')$ are never derivable or if some member of $\text{body}^-(r')$ has been derived by higher-ranked or unrelated rules. This is reflected by the identity of conditions IIa and 2a/b in all three approaches, respectively. Although this is

⁸Condition $\text{body}^-(r') \cap X \neq \emptyset$ in (4) is obsolete since $r' \notin \Gamma_{\Pi}^X$.

arguably a sensible strategy, it leads to the loss of preferred answer sets on programs like

$$\begin{array}{ll} r_1 : & a \leftarrow \text{not } b \\ r_2 : & b \leftarrow . \end{array} \quad r_2 < r_1$$

Let us now discuss the differences among the approaches. The difference between D- and w-preference can be directly read off Definition 1 and 4; it manifests itself in Condition IIb and leads to the following relationship.

Theorem 6 *Every D-preferred answer set is w-preferred.*

Example (3) shows that the converse does not hold.

Interestingly, a similar relationship is obtained between w- and B-preference. In fact, Definition 8 can be interpreted as a weakening of Definition 5 by dropping Condition 0 and weakening Condition 2 (via 2c). We thus obtain the following result.

Theorem 7 *Every w-preferred answer set is B-preferred.*

Example (5) shows that the converse does not hold.

We obtain the following summarizing result by letting $\mathcal{AS}(\Pi) = \{X \mid C_\Pi(X) = X\}$ and $\mathcal{AS}_P(\Pi, <) = \{X \in \mathcal{AS}(\Pi) \mid X \text{ is } P\text{-preferred}\}$ for $P = w, D, B$.

Theorem 8 *Let $(\Pi, <)$ be a statically ordered logic program. Then, we have:*

$$\mathcal{AS}_D(\Pi, <) \subseteq \mathcal{AS}_w(\Pi, <) \subseteq \mathcal{AS}_B(\Pi, <) \subseteq \mathcal{AS}(\Pi)$$

In principle, this hierarchy is induced by a decreasing interaction between groundedness and preference. While D-preference requires the full compatibility of both concepts, this interaction is already weakened in w-preference, before it is fully abandoned in B-preference. This is nicely reflected by the evolution of Condition 0 in definitions 2, 5, and 8.

Notably, groundedness as such is not the ultimate distinguishing factor, as demonstrated by the fact that prerequisite-free programs do not necessarily lead to the same preferred answer sets, as witnessed in (3) and (6). Rather it is the degree of interaction between groundedness and preferences that makes the difference.

8 Conclusion

The notion of preference seems to be pervasive in logic programming when it comes to knowledge representation. This is reflected by numerous approaches that aim at enhancing logic programming with preferences in order to improve knowledge representation capacities. Despite the large variety of approaches, however, only very little attention has been paid to their structural differences and sameness, finally leading to solid semantical underpinnings.

This work is a first step towards a systematic account to logic programming with preferences. We elaborated upon three different approaches that were originally defined in rather heterogenous ways. We obtained three alternative yet uniform ways of characterizing preferred answer sets (in terms of fixpoints, order preservation, and an axiomatic account). The underlying uniformity provided us with a deeper understanding of how and which answer sets are preferred in each approach. This has led to a clarification of their relationships and subtle differences. In particular, we revealed

that the investigated approaches yield an increasing number of answer sets depending on how tight they connect preference to groundedness.

An interesting technical result of this paper is given by the equivalences between the fixpoint operators and the standard logic programming operators applied to the correspondingly transformed programs (cf. Theorem 2 and 4). This opens the avenue for further concepts of preference handling on the basis of the alternating fixpoint theory and its issuing semantics. Further research includes dynamic preferences and more efficient algorithms for different semantics in a unifying way.

Acknowledgements. This work was supported by DFG under grant FOR 375/1-1, TP C.

References

- [Baader and Hollunder, 1993] F. Baader and B. Hollunder. How to prefer more specific defaults in terminological default logic. In *Proc. IJCAI'93*, p 669–674, 1993.
- [Brewka and Eiter, 1999] G. Brewka and T. Eiter. Preferred answer sets for extended logic programs. *Artificial Intelligence*, 109(1-2):297–356, 1999.
- [Brewka and Eiter, 2000] G. Brewka and T. Eiter. Prioritizing default logic. In St. Hölldobler, ed, *Intellectics and Computational Logic*. Kluwer, 2000. To appear.
- [Brewka, 1994] G. Brewka. Adding priorities and specificity to default logic. In L. Pereira and D. Pearce, eds, *Proc. JELIA'94*, p 247–260. Springer, 1994.
- [Delgrande et al., 2000] J. Delgrande, T. Schaub, and H. Tompits. Logic programs with compiled preferences. In *Proc. ECAI 2000*, p 392–398. IOS Press, 2000.
- [Gelfond and Lifschitz, 1991] M. Gelfond and V. Lifschitz. Classical negation in logic programs and deductive databases. *New Generation Computing*, 9:365–385, 1991.
- [Gelfond and Son, 1997] M. Gelfond and T. Son. Reasoning with prioritized defaults. In J. Dix, L. Pereira, and T. Przytycki, eds, *Workshop on Logic Programming and Knowledge Representation*, p 164–223. Springer, 1997.
- [Rintanen, 1995] J. Rintanen. On specificity in default logic. In *Proc. IJCAI'95*, p 1474–1479. Morgan Kaufmann, 1995.
- [Sakama and Inoue, 1996] C. Sakama and K. Inoue. Representing priorities in logic programs. In M. Maher, ed, *Proc. JCSLP'96*, p 82–96. MIT Press, 1996.
- [van Gelder, 1993] A. van Gelder. The alternating fixpoint of logic programs with negation. *J. Computer and System Science*, 47:185–120, 1993.
- [Wang et al., 2000] K. Wang, L. Zhou, and F. Lin. Alternating fixpoint theory for logic programs with priority. In *Proc. Int'l Conf. Computational Logic*, p 164–178. Springer, 2000.
- [Zhang and Foo, 1997] Y. Zhang and N. Foo. Answer sets for prioritized logic programs. In J. Maluszynski, ed, *Proc. ISLP'97*, p 69–84. MIT Press, 1997.