The Well-Founded Semantics
is a Stratified Fitting Semantics

Pascal Hitzler and Matthias Wendt

Knowledge Representation and Reasoning Group, Artificial Intelligence Institute
Department of Computer Science, Dresden University of Technology
Dresden, Germany
{phitzler,mw177754}@inf.tu-dresden.de
www.wv.inf.tu-dresden.de

Abstract

Part of the theory of logic programming and nonmonotonic reasoning concerns the study of fixed-point semantics for these paradigms. While several different semantics have been proposed, and some have been more successful than others, the exact relationships between the approaches have not yet been fully understood. In this paper, we give new characterizations, using level mappings, of the Fitting semantics, the well-founded semantics, and the weakly perfect model semantics. The results will unmask the well-founded semantics as a stratified version of the Fitting semantics.

Contents

1 Introduction

One of the very stimulating research questions in logic programming and nonmonotonic reasoning has been the search for an appropriate declarative understanding of negation. Several different semantics have been proposed, see e.g. [Sub99], each being more or less convincing, depending on one's point of view, which may be that of a programmer, or motivated by common-sense reasoning.

Among the semantics based on three-valued logics, the Fitting semantics [Fit85] and the well-founded semantics [VGRS91] are prominent and widely acknowledged choices. Theoretical relationships between them have been etablished, e.g. in [Fit0x] by using lattice-based logic programming in four-valued logic, an approach which was recently extended in [DMT00]. The development of the weakly perfect model semantics, due to [PP90], was motivated by the intuition that recursion should be allowed through positive information, but not through negation. As such, it was developed out of the notion of stratification [ABW88, Prz88]. We will see that the well-founded semantics achieves this very goal in a much better way than the weakly perfect model semantics. This will probably be no surprise for the specialist, although we are not aware of any formal argument for this case, and we believe that a presentation from our point of view is rather clear and convincing.

We will provide new, and uniform, characterizations, based on level mappings, of the Fitting semantics, the well-founded semantics, and the weakly perfect model semantics. Level mappings are mappings from Herbrand bases to ordinals, i.e. they induce orderings on the set of all ground atoms while disallowing infinite descending chains. They have been studied in termination analysis for logic programming, e.g. in [Bez89, AP93, Mar96], where they appear naturally, they have been used for defining classes of programs with desirable semantic properties, e.g. in [ABW88, Prz88, Cav91], they are intertwined with topological investigations of fixed-point semantics in logic programming, as studied e.g. in [Fit94, Sed95, Hit01, HS0x], and are relevant to some aspects of the study of relationships between logic programming and artificial neural networks [HKS99]. Our motivation, however, is quite different. We will employ level mappings in order to give uniform characterizations of different semantics in nonmonotonic reasoning, and we will this way employ them to unmask the well-founded semantics as a stratified version of the Fitting semantics, i.e. we will see that the difference between the Fitting semantics and the well-founded semantics is, in a nutshell, that at some particular point the former prevents recursion, while the latter at the same point prevents only recursion through negation.

The paper is structured as follows. Section 2 contains preliminaries which are needed to make the paper relatively self-contained. Section 3 contains our new characterization of the Fitting semantics, while Section 4 covers the new characterization of the well-founded semantics. In Section 5 the weakly perfect model semantics will be studied. We conclude with a summary and a discussion of further work in Section 6.

2 Preliminaries and Notation

A (*normal*) *logic program* is a finite set of (universally quantified) *clauses* of the form $\forall (A \leftarrow A_1 \wedge \cdots \wedge A_n \wedge \neg B_1 \wedge \cdots \wedge \neg B_m)$, commonly written as $A \leftarrow A_1, \ldots, A_n, \neg B_1, \ldots, \neg B_m$,

2

where A, A_i, and B_j, for $i = 1, \ldots, n$ and $j = 1, \ldots, m$, are atoms over some given first order language. A is called the *head* of the clause, while the remaining atoms make up the *body* of the clause, and depending on context, a body of a clause will be a set of literals (i.e. atoms or negated atoms) or the conjunction of these literals. Care will be taken that this identification does not cause confusion. We allow a body, i.e. a conjunction, to be empty, in which case it always evaluates to true. A clause with empty body is called a *unit clause* or a *fact*. A clause is called *definite*, if it contains no negation symbol. A program is called *definite* if it consists only of definite clauses. We will usually denote atoms with A or B, and literals, which may be atoms or negated atoms, by L or K.

Given a logic program P, we can extract from it the components of a first order language. The corresponding set of ground atoms, i.e. the *Herbrand base* of the program, will be denoted by B_P. For a subset $I \in B_P$, we set $\neg I = \{\neg A \mid A \in B_P\}$. The set of all ground instances of P with respect to B_P will be denoted by $\mathsf{ground}(P)$. A (*three-valued* or *partial*) *interpretation* I for P is a subset of $B_P \cup \neg B_P$ which is *consistent*, i.e. whenever $A \in I$ then $\neg A \notin I$. We say that A is *true with respect to* (or *in*) I if $A \in I$, we say that A is *false with respect to* (or *in*) I if $\neg A \in I$, and if neither is the case, we say that A is *undefined with respect to* (or *in*) I. A body, i.e. a conjunction of literals, is true in an interpretation I if every literal in the body is true in I, it is false in I if one of its literals is false in I, and otherwise it is undefined in I. For a negated literal $L = \neg A$ we will find it convenient to write $\neg L \in I$ if $A \in I$. By I_P we denote the set of all (three-valued) interpretations of P. It is a cpo via set-inclusion, i.e. it contains the empty set as least element, and every ascending chain has a supremum, namely its union. A *model* of P is an interpretation $I \in I_P$ such that for each clause $A \leftarrow \mathsf{body}$ we have that $\mathsf{body} \subseteq I$ implies $A \in I$. A *total* interpretation is an interpretation I such that no $A \in B_P$ is undefined in I.

For an interpretation I and a program P, an *I-partial level mapping* for P is a partial mapping $l : B_P \to \alpha$ with domain $\mathsf{dom}(l) = \{A \mid A \in I \text{ or } \neg A \in I\}$, where α is some (countable) ordinal. We extend every level mapping to literals by setting $l(\neg A) = l(A)$ for all $A \in \mathsf{dom}(l)$. A (*total*) *level mapping* is a total mapping $l : B_P \to \alpha$ for some (countable) ordinal α.

Given a normal logic program P and some $I \in I_P$, we say that $U \subseteq B_P$ is an *unfounded set* (*of P*) *with respect to I* if each atom $A \in U$ satisfies the following condition: For each clause $A \leftarrow \mathsf{body}$ in $\mathsf{ground}(P)$ (at least) one of the following holds.

(Ui) Some (positive or negative) literal in body is false in I.

(Uii) Some (non-negated) atom in body occurs in U.

Given a normal logic program P, we define the following operators on I_P. $T_P(I)$ is the set of all $A \in B_P$ such that there exists a clause $A \leftarrow \mathsf{body}$ in $\mathsf{ground}(P)$ such that body is true in I. $F_P(I)$ is the set of all $A \in B_P$ such that for all clauses $A \leftarrow \mathsf{body}$ in $\mathsf{ground}(P)$ we have that body is false in I. $U_P(I)$ is the greatest unfounded set (of P) with respect to I, which always exists due to [VGRS91]. Finally, define

$$\Phi_P(I) = T_P(I) \cup \neg F_P(I) \qquad \text{and}$$
$$W_P(I) = T_P(I) \cup \neg U_P(I)$$

3

for all $I \in I_P$. The operator Φ_P is due to [Fit85], while W_P is due to [VGRS91]. Both are monotonic on the cpo I_P, hence have a least fixed point by the Tarski fixed-point theorem, and we can obtain these fixed points by defining, for each monotonic operator F, that $F \uparrow 0 = \emptyset$, $F \uparrow (\alpha + 1) = F(F \uparrow \alpha)$ for any ordinal α, and $F \uparrow \beta = \bigcup_{\gamma < \beta} F \uparrow \gamma$ for any limit ordinal β, and the least fixed point of F is obtained as $F \uparrow \alpha$ for some ordinal α. The least fixed point of Φ_P is called the *Kripke-Kleene model* or *Fitting model* of P, determining the *Fitting semantics* of P, while the least fixed point of W_P is called the *well-founded model* of P, giving the *well-founded semantics* of P.

Given a program P, we define the operator T_P^+ on subsets of B_P by $T_P^+(I) = T_P(I \cup \neg(B_P \setminus I))$. It is well-known that for definite programs this operator is monotonic on the set of all subsets of B_P, with respect to subset inclusion. Indeed it is Scott-continuous [Llo88, SHLG94] and, via Kleene's fixed-point theorem, achieves its least fixed point M as the supremum of the iterates $T_P^+ \uparrow n$ for $n \in \mathbb{N}$. So M is *the least two-valued model* of P. In turn, we can identify M with the total interpretation $M \cup \neg(B_P \setminus M)$, which we will call the *definite (partial) model* of P.

In order to avoid confusion, we will use the following terminology: the notion of *interpretation* will by default denote consistent subsets of $B_P \cup \neg B_P$, i.e. interpretations in three-valued logic. We will sometimes emphasize this point by using the notion *partial interpretation*. By *two-valued interpretations* we mean subsets of B_P. Both interpretations and two-valued interpretations are ordered by subset inclusion. Each two-valued interpretation I can be identified with the partial interpretation $I' = I \cup \neg(B_P \setminus I)$. Note however, that in this case I' is always a maximal element in the ordering for partial interpretations, while I is in general not maximal as a two-valued interpretation.

3 Fitting Semantics

We give a new characterization, using level mappings, of the Fitting semantics.

3.1 Definition Let P be a normal logic program, I be a model of P, and l be an I-partial level mapping for P. We say that P *satisfies* (F) *with respect to I and l*, if each $A \in \mathsf{dom}(l)$ satisfies one of the following conditions.

(Fi) $A \in I$ and there exists a clause $A \leftarrow L_1, \ldots, L_n$ in $\mathsf{ground}(P)$ with head A such that $L_i \in I$ and $l(A) > l(L_i)$ for all i.

(Fii) $\neg A \in I$ and for each clause $A \leftarrow L_1, \ldots, L_n$ in $\mathsf{ground}(P)$ with head A there exists i with $\neg L_i \in I$ and $l(A) > l(L_i)$.

If $A \in \mathsf{dom}(l)$ satisfies (Fi), then we say that A *satisfies* (Fi) *with respect to I and l*, and similarly if $A \in \mathsf{dom}(l)$ satisfies (Fii).

3.2 Theorem Let P be a normal logic program with Fitting model M. Then M is the greatest model among all models I, for which there exists an I-partial level mapping l for P such that P satisfies (F) with respect to I and l.

Proof: Let M_P be the Fitting model of P and define the M_P-partial level mapping l_P as follows: $l_P(A) = \alpha$, where α is the least ordinal such that A is not undefined in $\Phi_P \uparrow (\alpha + 1)$. The proof will be established by showing the following facts: (1) P satisfies (F) with respect to M_P and l_P. (2) If I is a model of P and l an I-partial level mapping such that P satisfies (F) with respect to I and l, then $I \subseteq M_P$.

(1) Let $A \in \mathsf{dom}(l_P)$ and $l_P(A) = \alpha$. We consider two cases.

(Case i) If $A \in M_P$, then $A \in T_P(\Phi_P \uparrow \alpha)$, hence there exists a clause $A \leftarrow \mathtt{body}$ in $\mathsf{ground}(P)$ such that $\mathtt{body}$ is true in $\Phi_P \uparrow \alpha$. Thus, for all $L_i \in \mathtt{body}$ we have that $L_i \in \Phi_P \uparrow \alpha$, and hence $l_P(L_i) < \alpha$ and $L_i \in M_P$ for all i. Consequently, A satisfies (Fi) with respect to M_P and l_P.

(Case ii) If $\neg A \in M_P$, then $A \in F_P(\Phi_P \uparrow \alpha)$, hence for all clauses $A \leftarrow \mathtt{body}$ in $\mathsf{ground}(P)$ there exists $L \in \mathtt{body}$ with $\neg L \in \Phi_P \uparrow \alpha$ and $l_P(L) < \alpha$, hence $\neg L \in M_P$. Consequently, A satisfies (Fii) with respect to M_P and l_P, and we have established that fact (1) holds.

(2) We show via transfinite induction on $\alpha = l(A)$, that whenever $A \in I$ (respectively, $\neg A \in I$), then $A \in \Phi_P \uparrow (\alpha + 1)$ (respectively, $\neg A \in \Phi_P \uparrow (\alpha + 1)$). For the base case, note that if $l(A) = 0$, then $A \in I$ implies that A occurs as the head of a fact in $\mathsf{ground}(P)$, hence $A \in \Phi_P \uparrow 1$, and $\neg A \in I$ implies that there is no clause with head A in $\mathsf{ground}(P)$, hence $\neg A \in \Phi_P \uparrow 1$. So assume now that the induction hypothesis holds for all $B \in B_P$ with $l(B) < \alpha$. We consider two cases.

(Case i) If $A \in I$, then it satisfies (Fi) with respect to I and l. Hence there is a clause $A \leftarrow \mathtt{body}$ in $\mathsf{ground}(P)$ such that $\mathtt{body} \subseteq I$ and $l(K) < \alpha$ for all $K \in \mathtt{body}$. Hence $\mathtt{body} \subseteq M_P$ by induction hypothesis, and since M_P is a model of P we obtain $A \in M_P$.

(Case ii) If $\neg A \in I$, then A satisfies (Fii) with respect to I and l. Hence for all clauses $A \leftarrow \mathtt{body}$ in $\mathsf{ground}(P)$ we have that there is $K \in \mathtt{body}$ with $\neg K \in I$ and $l(K) < \alpha$. Hence for all these K we have $\neg K \in M_P$ by induction hypothesis, and consequently for all clauses $A \leftarrow \mathtt{body}$ in $\mathsf{ground}(P)$ we obtain that $\mathtt{body}$ is false in M_P. Since $M_P = \Phi_P(M_P)$ is a fixed point of the Φ_P-operator, we obtain $\neg A \in M_P$. This establishes fact (2) and concludes the proof. ∎

The following corollary follows immediately as a special case, and is closely related to results reported in [HS99, HS0x].

3.3 Corollary A normal logic program P has a total Fitting model if and only if there is a total model I of P and a (total) level mapping l for P such that P satisfies (F) with respect to I and l.

4 Well-Founded Semantics

We will provide a new characterization, via level mappings, of the well-founded semantics, and we will argue that from this new point of view the well-founded semantics can be understood as a stratified version of the Fitting semantics.

Let us first recall the definition of a (locally) stratified program, due to [ABW88, Prz88]: A normal logic program is called *locally stratified* if there exists a (total) level mapping $l : B_P \to \alpha$, for some ordinal α, such that for each clause $A \leftarrow A_1, \ldots, A_n, \neg B_1, \ldots, \neg B_m$ in $\mathsf{ground}(P)$ we have that $l(A) \geq l(A_i)$ and $l(A) > l(B_j)$ for all $i = 1, \ldots, n$ and $j = 1, \ldots, m$.

The notion of (locally) stratifed program was developed with the idea of preventing *recursion through negation*, while allowing recursion through positive dependencies. (Locally) stratified programs have total well-founded models.

There exist locally stratified programs which do not have a total Fitting semantics and vice versa. In fact, condition (Fii) requires a strict decrease of level between the head and a literal in the rule, independent of this literal being positive or negative. But, on the other hand, condition (Fii) imposes no further restrictions on the remaining body literals, while the notion of local stratification does. These considerations motivate the substitution of condition (Fii) by the condition (WFii), as given in the following definition.

4.1 Definition Let P be a normal logic program, I be a model of P, and l be an I-partial level mapping for P. We say that P *satisfies* (WF) *with respect to I and l*, if each $A \in \mathsf{dom}(l)$ satisfies one of the following conditions.

(WFi) $A \in I$ and there exists a clause $A \leftarrow L_1, \ldots, L_n$ in $\mathsf{ground}(P)$ with head A such that $L_i \in I$ and $l(A) > l(L_i)$ for all i.

(WFii) $\neg A \in I$ and for each clause $A \leftarrow A_1, \ldots, A_n, \neg B_1, \ldots, \neg B_m$ in $\mathsf{ground}(P)$ with head A (at least) one of the following conditions holds:

(WFiia) There exists $i \in \{1, \ldots, n\}$ with $\neg A_i \in I$ and $l(A) \geq l(A_i)$.

(WFiib) There exists $j \in \{1, \ldots, m\}$ with $B_j \in I$ and $l(A) > l(B_j)$.

If $A \in \mathsf{dom}(l)$ satisfies (WFi), then we say that A *satisfies* (WFi) *with respect to I and l*, and similarly if $A \in \mathsf{dom}(l)$ satisfies (WFii).

We note that conditions (Fi) and (WFi) are identical. Indeed, replacing (WFi) by a stratified version such as the following is not satisfactory.

(SFi) $A \in I$ and there exists a clause $A \leftarrow A_1, \ldots, A_n, \neg B_1, \ldots, \neg B_m$ in $\mathsf{ground}(P)$ with head A such that $A_i, B_j \in I$, $l(A) \geq l(A_i)$, and $l(A) > l(B_j)$ for all i and j.

If we replace condition (WFi) by condition (SFi), then it is not guaranteed that for any given program there is a maximal model satisfying the desired properties: Consider the program consisting of the two clauses $p \leftarrow p$ and $q \leftarrow \neg p$, and the two (total) models $\{p, \neg q\}$ and $\{\neg p, q\}$, which are incomparable, and the level mapping l with $l(p) = 0$ and $l(q) = 1$.

So, in the light of Theorem 3.2, Definition 4.1 should provide a natural "stratified version" of the Fitting semantics. And indeed it does, and furthermore, the resulting semantics coincides with the well-founded semantics, which is a very satisfactory result.

4.2 Theorem Let P be a normal logic program with well-founded model M. Then M is the greatest model among all models I, for which there exists an I-partial level mapping l for P such that P satisfies (WF) with respect to I and l.

Proof: Let M_P be the well-founded model of P and define the M_P-partial level mapping l_P as follows: $l_P(A) = \alpha$, where α is the least ordinal such that A is not undefined in $W_P \uparrow (\alpha + 1)$. The proof will be established by showing the following facts: (1) P satisfies (WF) with respect

to M_P and l_P. (2) If I is a model of P and l an I-partial level mapping such that P satisfies (WF) with respect to I and l, then $I \subseteq M_P$.

(1) Let $A \in \mathsf{dom}(l_P)$ and $l_P(A) = \alpha$. We consider two cases.

(Case i) If $A \in M_P$, then $A \in T_P(W_P \uparrow \alpha)$, hence there exists a clause $A \leftarrow \mathtt{body}$ in $\mathsf{ground}(P)$ such that $\mathtt{body}$ is true in $W_P \uparrow \alpha$. Thus, for all $L_i \in \mathtt{body}$ we have that $L_i \in W_P \uparrow \alpha$, and hence $l_P(L_i) < \alpha$ and $L_i \in M_P$ for all i. Consequently, A satisfies (WFi) with respect to M_P and l_P.

(Case ii) If $\neg A \in M_P$, then $A \in U_P(W_P \uparrow \alpha)$, i.e. A is contained in the greatest unfounded set of P with respect to $W_P \uparrow \alpha$. Hence for each clause $A \leftarrow \mathtt{body}$ in $\mathsf{ground}(P)$, at least one of (Ui) or (Uii) holds with respect to $W_P \uparrow \alpha$ and the unfounded set $U_P(W_P \uparrow \alpha)$. If (Ui) holds, then there exists some literal $L \in \mathtt{body}$ with $\neg L \in W_P \uparrow \alpha$, hence $l_P(L) < \alpha$ and condition (WFiia) (if L is an atom), respectively condition (WFiib) (if L is a negated atom), is satisfied by A with respect to M_P and l_P. If (Uii) holds, then some (non-negated) atom B in $\mathtt{body}$ occurs in $U_P(W_P \uparrow \alpha)$. Hence $l_P(B) \leq l_P(A)$ and A satisfies (WFiia) with respect to M_P and l_P. Thus we have established that fact (1) holds.

(2) We show via transfinite induction on $\alpha = l(A)$, that whenever $A \in I$ (respectively, $\neg A \in I$), then $A \in W_P \uparrow (\alpha + 1)$ (respectively, $\neg A \in W_P \uparrow (\alpha + 1)$). For the base case, note that if $l(A) = 0$, then $A \in I$ implies that A occurs as the head of a fact in $\mathsf{ground}(P)$, hence $A \in W_P \uparrow 1$. If $\neg A \in I$, then consider the set U of all atoms B with $l(B) = 0$ and $\neg B \in I$. We show that U is an unfounded set of P with respect to $I = M_P$, and this suffices since it implies $\neg A \in M_P$ by $A \in U$ and the fact that M_P is a fixed point of W_P. So let $C \in U$ and let $C \leftarrow \mathtt{body}$ be a clause in $\mathsf{ground}(P)$. Since $\neg C \in I$, and $l(C) = 0$, we have that C satisfies (WFiia) with respect to I and l, so condition (Uii) is satisfied and we have that U is an unfounded set of P with respect to I. Assume now that the induction hypothesis holds for all $B \in B_P$ with $l(B) < \alpha$. We consider two cases.

(Case i) If $A \in I$, then it satisfies (WFi) with respect to I and l. Hence there is a clause $A \leftarrow \mathtt{body}$ in $\mathsf{ground}(P)$ such that $\mathtt{body} \subseteq I$ and $l(K) < \alpha$ for all $K \in \mathtt{body}$. Hence $\mathtt{body} \subseteq W_P \uparrow \alpha$, and we obtain $A \in T_P(W_P \uparrow \alpha)$ as required.

(Case ii) If $\neg A \in I$, consider the set U of all atoms B with $l(B) = \alpha$ and $\neg B \in I$. We show that U is an unfounded set of P with respect to M_P, and this suffices since it implies $\neg A \in M_P$ by $A \in U$ and the fact that M_P is a fixed point of W_P. So let $C \in U$ and let $C \leftarrow \mathtt{body}$ be a clause in $\mathsf{ground}(P)$. Since $\neg C \in I$, we have that C satisfies (WFii) with respect to I and l. If there is a literal $L \in \mathtt{body}$ with $\neg L \in I$ and $l(L) < l(C)$, then by induction hypothesis we obtain $\neg L \in M_P$, so condition (Ui) is satisfied for the clause $C \leftarrow \mathtt{body}$ with respect to M_P and U. In the remaining case we have that C satisfies condition (WFiia), and there exists an atom $B \in \mathtt{body}$ with $\neg B \in I$ and $l(B) = l(C)$, hence $B \in U$, i.e. condition (Uii) is satisfied for the clause $C \leftarrow \mathtt{body}$ with respect to M_P and U. Hence U is an unfounded set of P with respect to M_P. ∎

As a special case, we immediately obtain the following corollary, which was obtained directly in [Wen02].

4.3 Corollary A normal logic program P has a total well-founded model if and only if there is a total model I of P and a (total) level mapping l such that P satisfies (WF) with respect to I and l.

5 Weakly Perfect Model Semantics

We have obtained new characterizations of the Fitting semantics and the well-founded semantics, and argued that the well-founded semantics is a stratified version of the Fitting semantics. Our argumentation is based on the key intuition underlying the notion of stratification, that recursion should be allowed through positive dependencies, but be forbidden through negative dependencies. As we have seen in Theorem 4.2, the well-founded semantics provides this for a setting in three-valued logic. Historically, a different semantics, given by the so-called weakly perfect model associated with each program, was proposed in [PP90] in order to carry over the intuition underlying the notion of stratification to a three-valued setting. In the following, we will characterize weakly perfect models via level mappings, in the spirit of Theorems 3.2 and 4.2. We will thus have obtained uniform characterizations of the Fitting semantics, the well-founded semantics, and the weakly perfect model semantics, which makes it possible to easily compare them.

5.1 Definition Let P be a normal logic program, I be a model of P and l be an I-partial level mapping for P. We say that P *satisfies* (WS) *with respect to I and l*, if each $A \in \mathsf{dom}(l)$ satisfies one of the following conditions.

(WSi) $A \in I$ and there exists a clause $A \leftarrow L_1, \ldots, L_n \in \mathsf{ground}(P)$ such that $L_i \in I$ and $l(A) > l(L_i)$ for all $i = 1, \ldots, n$.

(WSii) $\neg A \in I$ and for each clause $A \leftarrow A_1, \ldots, A_n, \neg B_1, \ldots, \neg B_m \in \mathsf{ground}(P)$ (at least) one of the following three conditions holds.

    (WSiia) There exists i such that $\neg A_i \in I$ and $l(A) > l(A_i)$.

    (WSiib) For all k we have $l(A) \geq l(A_k)$, for all j we have $l(A) > l(B_j)$, and there exists i with $\neg A_i \in I$.

    (WSiic) There exists j such that $B_j \in I$ and $l(A) > l(B_j)$.

We observe that the condition (WSii) in the above theorem is more general than (Fii), and more restrictive than (WFii).

We will see below in Theorem 5.3, that Definition 5.1 captures the weakly perfect model, in the same way in which Definitions 3.1 and 4.1 capture the Fitting model, respectively the well-founded model.

In order to proceed with this, we first need to recall the definition of weakly perfect models due to [PP90], and we will do this next. For ease of notation, it will be convenient to consider (countably infinite) propositional programs instead of programs over a first-order language. This is both common practice and no restriction, because the ground instantiation $\mathsf{ground}(P)$ of a given program P can be understood as a propositional program which may consist of a countably infinite number of clauses. Let us remark that our definition below differs slightly from that given in [PP90], and we will return to this point later. It nevertheless leads to exactly the same notion of weakly stratified program.

Let P be a (countably infinite propositional) normal logic program. An atom $A \in B_P$ *refers to* an atom $B \in B_P$ if B or $\neg B$ occurs as a body literal in a clause $A \leftarrow \mathtt{body}$ in P. A *refers negatively to* B if $\neg B$ occurs as a body literal in such a clause. We say that A *depends*

on B if the pair (A, B) is in the transitive closure of the relation *refers to*, and we write this as $B \leq A$. We say that A *depends negatively on* B if there are $C, D \in B_P$ such that C refers negatively to D and one of the following holds: (1) $C \leq A$ or $C = A$ (the latter meaning identity). (2) $B \leq D$ or $B = D$. We write $B < A$ in this case. For $A, B \in B_P$, we write $A \sim B$ if either $A = B$, or A and B depend negatively on each other, i.e. if $A < B$ and $B < A$ hold. The relation $\sim$ is an equivalence relation and its equivalence classes are called *components* of P. A component is *trivial* if it consists of a single element A with $A \not< A$.

Let C_1 and C_2 be two components of a program P. We write $C_1 \prec C_2$ if and only if $C_1 \neq C_2$ and for all $A_1 \in C_1$ there is $A_2 \in C_2$ with $A_1 < A_2$. A component C_1 is called *minimal* if there is no component C_2 with $C_2 \prec C_1$.

Given a normal logic program P, the *bottom stratum* $S(P)$ of P is the union of all minimal components of P. The *bottom layer* of P is the subprogram $L(P)$ of P which consists of all clauses from P with heads belonging to $S(P)$.

Given a (partial) interpretation I of P, we define the *reduct of P with respect to I* as the program P/I obtained from P by performing the following reductions. (1) Remove from P all clauses which contain a body literal L such that $\neg L \in I$ or whose head belongs to I. (2) Remove from all remaining clauses all body literals L with $L \in I$. (3) Remove from the resulting program all non-unit clauses, whose heads appear also as unit clauses in the program.

5.2 Definition The *weakly perfect model* M_P of a program P is defined by transfinite induction as follows. Let $P_0 = P$ and $M_0 = \emptyset$. For each (countable) ordinal $\alpha > 0$ such that programs P_δ and partial interpretations M_δ have already been defined for all $\delta < \alpha$, let

$$N_\alpha = \bigcup_{0 < \delta < \alpha} M_\delta,$$

$$P_\alpha = P/N_\alpha,$$

R_α is the set of all atoms which are undefined in N_α

and were eliminated from P by reducing it with respect to N_α,

$$S_\alpha = S\left(P_\alpha\right), \text{ and}$$

$$L_\alpha = L\left(P_\alpha\right).$$

The construction then proceeds with one of the following three cases. (1) If P_α is empty, then the construction stops and $M_P = N_\alpha \cup \neg R_\alpha$ is the (*total*) *weakly perfect model* of P. (2) If the bottom stratum S_α is empty or if the bottom layer L_α contains a negative literal, then the construction also stops and $M_P = N_\alpha \cup \neg R_\alpha$ is the (*partial*) *weakly perfect model* of P. (3) In the remaining case L_α is a definite program, and we define $M_\alpha = H \cup \neg R_\alpha$, where H is the definite (partial) model of L_α, and the construction continues.

For every α, the set $S_\alpha \cup R_\alpha$ is called the *α-th stratum* of P and the program L_α is called the *α-th layer* of P.

A *weakly stratified program* is a program with a total weakly perfect model. The set of its strata is then called its *weak stratification*.

Let us return to the remark made earlier that our definition of weakly perfect model, as given in Definition 5.2, differs slightly from the version introduced in [PP90]. In order to obtain

the original definition, points (2) and (3) of Definition 5.2 have to be replaced as follows: (2) If the bottom stratum S_α is empty or if the bottom layer L_α *has no least two-valued model*, then the construction stops and $M_P = N_\alpha \cup \neg R_\alpha$ is the (partial) weakly perfect model of P. (3) In the remaining case L_α *has a least two-valued model*, and we define $M_\alpha = H \cup \neg R_\alpha$, where H is the partial model of L_α corresponding to its least two-valued model, and the construction continues.

The original definition is more general due to the fact that every definite program has a least two-valued model. However, while the least two-valued model of a definite program can be obtained as the least fixed point of the monotonic (and even Scott-continuous) operator T_P^+, we know of no similar result, or general operator, for obtaining the least two-valued model, if existent, of progams which are not definite. The original definition therefore seems to be rather awkward, and indeed, even in [PP90], when defining weakly stratified programs, the more general version was dropped in favour of requiring definite layers. So Definition 5.2 is an adaptation taking the original notion of weakly stratified program into account, and appears to be more natural. In the following, the notion of *weakly perfect model* will refer to Definition 5.2.

5.3 Theorem Let P be a normal logic program with weakly perfect model M_P. Then M_P is the greatest model among all models I, for which there exists an I-partial level mapping l for P such that P satisfies (WS) with respect to I and l.

We prepare the proof of Theorem 5.3 by introducing some notation, which will make the presentation much more transparent.

It will be very convenient to consider level mappings which map into *pairs* (β, n) of ordinals, where $n \leq \omega$, the least infinite ordinal. So let α be a (countable) ordinal and consider the set $\mathcal{A}$ of all pairs (β, n), where $\beta < \alpha$ and $n \leq \omega$. Of course $\mathcal{A}$ endowed with the lexicographic ordering is isomorphic to an ordinal. So any mapping from B_P to $\mathcal{A}$ can be considered to be a level mapping.

Let P be a normal logic program with (partial) weakly perfect model M_P. Then define the M_P-partial level mapping l_P as follows: $l_P(A) = (\beta, n)$, where $A \in S_\beta \cup R_\beta$ and n is least with $A \in T_{L_\beta}^+ \uparrow (n+1)$, if such an n exists, and $n = \omega$ otherwise. We observe that if $l_P(A) = l_P(B)$ then there exists α with $A, B \in S_\alpha \cup R_\alpha$, and if $A \in S_\alpha \cup R_\alpha$ and $B \in S_\beta \cup R_\beta$ with $\alpha < \beta$, then $l(A) < l(B)$.

We next define model-consistent subsumption due to [Wen02].

5.4 Definition Let P and Q be two programs and let I be an interpretation.

1. If $C_1 = (A \leftarrow L_1, \ldots, L_m)$ and $C_2 = (B \leftarrow K_1, \ldots, K_n)$ are two clauses, then we say that C_1 *subsumes* C_2, written $C_1 \preccurlyeq C_2$, if $A = B$ and $\{L_1, \ldots, L_m\} \subseteq \{K_1, \ldots, K_n\}$.

2. We say that P *subsumes* Q, written $P \preccurlyeq Q$, if for each clause C_1 in P there exists a clause C_2 in Q with $C_1 \preccurlyeq C_2$.

3. We say that P *subsumes* Q *model-consistently* (*with respect to I*), written $P \preccurlyeq_I Q$, if the following conditions hold. (i) For each clause $C_1 = (A \leftarrow L_1, \ldots, L_m)$ in P there exists a clause $C_2 = (B \leftarrow K_1, \ldots, K_n)$ in Q with $C_1 \preccurlyeq C_2$ and $(\{K_1, \ldots, K_n\} \setminus \{L_1, \ldots, L_m\}) \subseteq$

I. (ii) For each clause $C_2 = (B \leftarrow K_1, \ldots, K_n)$ in Q with $\{K_1, \ldots, K_n\} \in I$ and $B \notin I$ there exists a clause C_1 in P such that $C_1 \preccurlyeq C_2$.

Definition 5.4 facilitates the proof of Theorem 5.3 by employing the following lemma.

5.5 Lemma With notation from Definiton 5.2, we have $P/N_\alpha \preccurlyeq_{N_\alpha} P$ for all α.

Proof: Condition 3(i) of Definition 5.4 holds because every clause in P/N_α is obtained from a clause in P by deleting body literals which are contained in N_α. Condition 3(ii) holds because for each clause in P with head $A \notin N_\alpha$ whose body is true under N_α, we have that $A \leftarrow$ is a fact in P/N_α. ∎

The next lemma establishes the induction step in part (2) of the proof of Theorem 5.3.

5.6 Lemma If I is a non-empty model of a (infinite propositional normal) logic program P' and l an I-partial level mapping such that P' satisfies (WS) with respect to I and l, then the following hold for $P = P'/\emptyset$.

(a) The bottom stratum $S(P)$ of P is non-empty and consists of trivial components only.

(b) The bottom layer $L(P)$ of P is definite.

(c) The definite (partial) model N of $L(P)$ is consistent with I in the following sense: we have $I' \subseteq N$, where I' is the restriction of I to all atoms which are not undefined in N.

(d) P/N satisfies (WS) with respect to $I \setminus N$ and l/N, where l/N is the restriction of l to the atoms in $I \setminus N$.

Proof: (a) Assume there exists some component $C \subseteq S(P)$ which is not trivial. Then there must exist atoms $A, B \in C$ with $A < B$, $B < A$, and $A \neq B$. Without loss of generality, we can assume that A is chosen such that $l(A)$ is minimal. Now let A' be any atom occuring in a clause with head A. Then $A > B > A \geq A'$, hence $A > A'$, and by minimality of the component we must also have $A' > A$, and we obtain that all atoms occuring in clauses with head A must be contained in C. We consider two cases.

(Case i) If $A \in I$, then there must be a fact $A \leftarrow$ in P, since otherwise by (WSi) we had a clause $A \leftarrow L_1, \ldots, L_n$ (for some $n \geq 1$) with $L_1, \ldots, L_n \in I$ and $l(A) > l(L_i)$ for all i, contradicting the minimality of $l(A)$. Since $P = P'/\emptyset$ we obtain that $A \leftarrow$ is the only clause in P with head A, contradicting the existence of $B \neq A$ with $B < A$.

(Case ii) If $\neg A \in I$, and since A was chosen minimal with respect to l, we obtain that condition (WSiib) must hold for each clause $A \leftarrow A_1, \ldots, A_n, \neg B_1, \ldots, \neg B_m$ with respect to I and l, and that $m = 0$. Furthermore, all A_i must be contained in C, as already noted above, and $l(A) \geq l(A_i)$ for all i by (WSiib). Also from (Case i) we obtain that no A_i can be contained in I. We have now established that for all A_i in the body of any clause with head A, we have $l(A) = l(A_i)$ and $\neg A_i \in I$. The same argument holds for all clauses with head A_i, for all i, and the argument repeats. Now from $A > B$ we obtain that there are $D, E \in C$ with $A \geq E$ (or $A = E$), $D \geq B$ (or $D = B$), and E refers negatively to D. As we have just seen, we obtain $\neg E \in I$ and $l(E) = l(A)$. Since E refers negatively to D, there is a clause with head E and $\neg D$ contained in the body of this clause. Since (WSii) holds for this clause,

there must be a literal L in the body with level less than $l(E)$, hence $l(L) < l(A)$ and $L \in C$ which is a contradiction. We thus have established that all components are trivial.

We show next that the bottom stratum is non-empty. Indeed, let A be an atom such that $l(A)$ is minimal. We will show that $\{A\}$ is a component. So assume it is not, i.e. that there is B with $B < A$. Then there exist $D_1, \dots, D_k$, for some $k \in \mathbb{N}$, such that $D_1 = A$, D_j refers to D_{j+1} for all $j = 1, \dots, k-1$, and D_k refers negatively to some B' with $B' \geq B$ (or $B' = B$).

We show next by induction that for all $j = 1, \dots, k$ the following statements hold: $\neg D_j \in I$, $B < D_j$, and $l(D_j) = l(A)$. Indeed note that for $j = 1$, i.e. $D_j = A$, we have that $B < D_j = A$ and $l(D_j) = l(A)$. Assuming $A \in I$, we obtain by minimality of $l(A)$ that $A \leftarrow$ is the only clause in $P = P'/\emptyset$ with head A, contradicting the existence of $B < A$. So $\neg A \in I$, and the assertion holds for $j = 1$. Now assume the assertion holds some $j < k$. Then obviously $D_{j+1} > B$. By $\neg D_j \in I$ and $l(D_j) = l(A)$, we obtain that (WSii) must hold, and by the minimality of $l(A)$ we infer that (WSiib) must hold and that no clause with head D_j contains negated atoms. So $l(D_{j+1}) = l(D_j) = l(A)$ holds by (WSiib) and minimality of $l(A)$. Furthermore, the assumption $D_{j+1} \in I$ can be rejected by the same argument as for A above, because then $D_{j+1} \leftarrow$ would be the only clause with head D_{j+1}, by minimality of $l(D_{j+1}) = l(A)$, contradicting $B < D_{j+1}$. This concludes the inductive proof.

Summarizing, we obtain that D_k refers negatively to B', and that $\neg D_k \in I$. But then there is a clause with head D_k and $\neg B'$ in its body which satisfies (WSii), contradicting the minimality of $l(D_k) = l(A)$. This concludes the proof of statement (a).

(b) According to [PP90] we have that whenever all components are trivial, then the bottom layer is definite. So the assertion follows from (a).

(c) Let $A \in I'$ be an atom with $A \notin N$, and assume without loss of generality that A is chosen such that $l(A)$ is minimal with these properties. Then there must be a clause $A \leftarrow \texttt{body}$ in P such that all literals in $\texttt{body}$ are true with respect to I', hence with respect to N by minimality of $l(A)$. Thus $\texttt{body}$ is true in N, and since N is a model of $L(P)$ we obtain $A \in N$, which contradicts our assumption.

Now let $A \in N$ be an atom with $A \notin I'$, and assume without loss of generality that A is chosen such that n is minimal with $A \in T^+_{L(P)} \uparrow (n+1)$. But then there is a definite clause $A \leftarrow \texttt{body}$ in $L(P)$ such that all atoms in $\texttt{body}$ are true with respect to $T^+_{L(P)} \uparrow n$, hence also with respect to I', and since I' is a model of $L(P)$ we obtain $A \in I'$, which contradicts our assumption.

Finally, let $\neg A \in I'$. Then we cannot have $A \in N$ since this implies $A \in I'$. So $\neg A \in N$ since N is a total model of $L(P)$.

(d) From Lemma 5.5, we know that $P/N \preccurlyeq_N P$. We distinguish two cases.

(Case i) If $I \setminus N \models A$, then there must exist a clause $A \leftarrow L_1, \dots, L_k$ in P such that $L_i \in I$ and $l(A) > l(L_i)$ for all i. Since it is not possible that $A \in N$, there must also be a clause in P/N which subsumes $A \leftarrow L_1, \dots, L_k$, and which therefore satisfies (WSi). So A satisfies (WSi).

(Case ii) If $\neg A \in I \setminus N$, then for each clause $A \leftarrow \texttt{body1}$ in P/N there must be a clause $A \leftarrow \texttt{body}$ in P which is subsumed by the former, and since $\neg A \in I$, we obtain that condition (WSii) must be satisfied by A, and by the clause $A \leftarrow \texttt{body}$. Since reduction with respect to N removes only body literals which are true in N, condition (WSii) is still met. ∎

We can now proceed with the proof.

Proof of Theorem 5.3: The proof will be established by showing the following facts: (1) P satisfies (WS) with respect to M_P and l_P. (2) If I is a model of P and l an I-partial level mapping such that P satisfies (WS) with respect to I and l, then $I \subseteq M_P$.

(1) Let $A \in \mathsf{dom}(l_P)$ and $l_P(A) = (\alpha, n)$. We consider two cases.

(Case i) If $A \in M_P$, then $A \in T_{L_\alpha}^+ \uparrow (n+1)$. Hence there exists a definite clause $A \leftarrow A_1, \ldots, A_k$ in L_α with $A_1, \ldots, A_k \in T_{L_\alpha}^+ \uparrow n$, so $A_1, \ldots, A_k \in M_P$ with $l_P(A) > l_P(A_i)$ for all i. Since $P/N_\alpha \preccurlyeq_{N_\alpha} P$ by Lemma 5.5, there must exist a clause $A \leftarrow A_1, \ldots, A_k, L_1, \ldots, L_m$ in P with literals $L_1, \ldots, L_m \in N_\alpha \subseteq M_P$, and we obtain $l_P(L_j) < l_P(A)$ for all $j = 1, \ldots, m$. So (WSi) holds in this case.

(Case ii) If $\neg A \in M_P$, then let $A \leftarrow A_1, \ldots, A_k, \neg B_1, \ldots, \neg B_m$ be a clause in P, noting that (WSii) is trivially satisfied in case no such clause exists. We consider the following two subcases.

(Subcase ii.a) Assume A is undefined in N_α and was eliminated from P by reducing it with respect to N_α, i.e. $A \in R_\alpha$. Then, in particular, there must be some $\neg A_i \in N_\alpha$ or some $B_j \in N_\alpha$, which yields $l_P(A_i) < l_P(A)$, respectively $l_P(B_j) < l_P(A)$, and hence one of (WSiia), (WSiic) holds.

(Subcase ii.b) Assume $\neg A \in H$, where H is the definite (partial) model of L_α. Since P/N_α subsumes P model-consistently with respect to N_α, we obtain that there must be some A_i with $\neg A_i \in H$, and by definition of l_P we obtain $l_P(A) = l_P(A_i) = (\alpha, \omega)$, and hence also $l_P(A_{i'}) \leq l_P(A_i)$ for all $i' \neq i$. Furthermore, since P/N_α is definite, we obtain that $\neg B_j \in N_\alpha$ for all j, hence $l_P(B_j) < l_P(A)$ for all j. So condition (WSiib) is satisfied.

(2) First note that for all models M, N of P with $M \subseteq N$ we have $(P/M)/N = P/(M \cup N) = P/N$ and $(P/N)/\emptyset = P/N$.

Let I_α denote I restricted to the atoms which are not undefined in $N_\alpha \cup R_\alpha$. It suffices to show the following: For all $\alpha > 0$ we have $I_\alpha \subseteq N_\alpha \cup R_\alpha$, and $I \setminus M_P = \emptyset$.

We next show by induction that if $\alpha > 0$ is an ordinal, then the following statements hold. (a) The bottom stratum of P/N_α is non-empty and consists of trivial components only. (b) The bottom layer of P/N_α is definite. (c) $I_\alpha \subseteq N_\alpha \cup R_\alpha$. (d) $P/N_{\alpha+1}$ satisfies (WS) with respect to $I \setminus N_{\alpha+1}$ and $l/N_{\alpha+1}$.

Note first that P satisfies the hypothesis of Lemma 5.6, hence also its consequences. So $P/N_1 = P/\emptyset$ satisfies (WS) with respect to $I \setminus N_1$ and l/N_1, and by application of Lemma 5.6 we obtain that statements (a) and (b) hold. For (c), note that no atom in R_1 can be true in I, because no atom in R_1 can appear as head of a clause in P, and apply Lemma 5.6 (c). For (d), apply Lemma 5.6, noting that $P/N_2 \preccurlyeq_{N_2} P$.

For α being a limit ordinal, we can show exactly as in the proof of Lemma 5.6 (d), that P satisfies (WS) with respect to $I \setminus N_\alpha$ and l/N_α. So Lemma 5.6 is applicable and statements (a) and (b) follow. For (c), let $A \in R_\alpha$. Then every clause in P with head A contains a body literal which is false in N_α. By induction hypothesis, this implies that no clause with head A in P can have a body which is true in I. So $A \notin I$. Together with Lemma 5.6 (c), this proves statement (c). For (d), apply again Lemma 5.6 (d), noting that $P/N_{\alpha+1} \preccurlyeq_{N_{\alpha+1}} P$.

For $\alpha = \beta + 1$ being a successor ordinal, we obtain by induction hypothesis that P/N_β satisfies the hypothesis of Lemma 5.6, so again statements (a) and (b) follow immediately from this lemma, and (c), (d) follow as in the case for α being a limit ordinal.

It remains to show that $I \setminus M_P = \emptyset$. Indeed by the transfinite induction argument just given we obtain that P/M_P satisfies (WS) with respect to $I \setminus M_P$ and l/M_P. If $I \setminus M_P$ is non-empty, then by Lemma 5.6 the bottom stratum $S(P/M_P)$ is non-empty and the bottom layer $L(P/M_P)$ is definite with definite (partial) model M. Hence by definition of the weakly perfect model M_P of P we must have that $M \subseteq M_P$ which contradicts the fact that M is the definite model of $L(P/M_P)$. Hence $I \setminus M_P$ must be empty which concludes the proof. ∎

We obtain the following corollary, previously reported, and proven directly, in [Wen02].

5.7 Corollary A normal logic program P is weakly stratified, i.e. has a total weakly perfect model, if and only if there is a total model I of P and a (total) level mapping l for P such that P satisfies (WS) with respect to I and l.

We also obtain the following corollary as a trivial consequence of our uniform characterizations by level mappings.

5.8 Corollary Let P be a normal logic progam with Fitting model M_{F}, weakly perfect model M_{WP}, and well-founded model M_{WF}. Then $M_{\mathrm{F}} \subseteq M_{\mathrm{WP}} \subseteq M_{\mathrm{WF}}$.

6   Conclusions and Further Work

We have obtained new characterizations of the Fitting semantics, the well-founded semantics, and the weakly perfect model semantics, and argued that the well-founded semantics is a stratified version of the Fitting semantics. Considering that the main motivation for the introduction of (weak) stratification was to restrict recursion through negation, we notice by comparing (WFii) and (WSii) that the well-founded semantics provides a much cleaner and more convincing way of achieving this.

Our approach, using level mappings, provides a way of comparing different semantics which is an alternative to the approach taken e.g. in [Fit0x, DMT00]. It provides uniform characterizations, and we believe that it should be applicable to most fixed-point semantics based on monotonic operators. In particular, it should be possible to employ our methods in order to characterize different forms of well-founded semantics for (extended) disjunctive logic programs, see [LMR92, BG94, LRS97], and also to the study of fixed-point semantics of logic programming in algebraic domains, as put forward in [RZ01, Hit02].

Under characterizations with level mappings, as proposed in this paper, a model should be computationally tractable, in a sense which remains to be specified by further research, if the corresponding partial level mapping maps into ω, the first infinite ordinal. This is certainly the case for Datalog, and it remains to be seen whether our results can be exploited for more efficient computation of the well-founded model, as in the currently evolving paradigm of answer set programming, see [Lif99, MT99, WZ00, SNS0x].

References

[ABW88]   Krzysztof R. Apt, Howard A. Blair, and Adrian Walker. Towards a theory of declarative knowledge. In Jack Minker, editor, *Foundations of Deductive Databases and Logic Programming*, pages 89–148. Morgan Kaufmann, Los Altos, CA, 1988.

[AP93]   Krzysztof R. Apt and Dino Pedreschi. Reasoning about termination of pure prolog programs. *Information and Computation*, 106:109–157, 1993.

[Bez89]   Marc Bezem. Characterizing termination of logic programs with level mappings. In Ewing L. Lusk and Ross A. Overbeek, editors, *Proceedings of the North American Conference on Logic Programming*, pages 69–80. MIT Press, Cambridge, MA, 1989.

[BG94]   Chitta Baral and Michael Gelfond. Logic programming and knowledge representation. *Journal of Logic Programming*, 19–20:73–148, 1994.

[Cav91]   Lawrence Cavedon. Acyclic programs and the completeness of SLDNF-resolution. *Theoretical Computer Science*, 86:81–92, 1991.

[DMT00]   Marc Denecker, V. Wiktor Marek, and Miroslaw Truszynski. Approximating operators, stable operators, well-founded fixpoints and applications in non-monotonic reasoning. In Jack Minker, editor, *Logic-based Artificial Intelligence*, chapter 6, pages 127–144. Kluwer Academic Publishers, Boston, 2000.

[Fit85]   Melvin Fitting. A Kripke-Kleene-semantics for general logic programs. *Journal of Logic Programming*, 2:295–312, 1985.

[Fit94]   Melvin Fitting. Metric methods: Three examples and a theorem. *Journal of Logic Programming*, 21(3):113–127, 1994.

[Fit0x]   Melvin Fitting. Fixpoint semantics for logic programming — A survey. *Theoretical Computer Science*, 200x. To appear.

[Hit01]   Pascal Hitzler. *Generalized Metrics and Topology in Logic Programming Semantics*. PhD thesis, Department of Mathematics, National University of Ireland, University College Cork, 2001.

[Hit02]   Pascal Hitzler. Resolution and logic programming in algebraic domains: Negation and defaults. Technical Report WV-02-05, Knowledge Representation and Reasoning Group, Department of Computer Science, Dresden University of Technology, Dresden, Germany, 2002.

[HKS99]   Steffen Hölldobler, Yvonne Kalinke, and Hans-Peter Störr. Approximating the semantics of logic programs by recurrent neural networks. *Applied Intelligence*, 11:45–58, 1999.

[HS99]   Pascal Hitzler and Anthony K. Seda. Characterizations of classes of programs by three-valued operators. In Michael Gelfond, Nicola Leone, and Gerald Pfeifer, editors, *Logic Programming and Nonmonotonic Reasoning, Proceedings of the 5th*

International Conference on Logic Programming and Non-Monotonic Reasoning
(LPNMR'99), El Paso, Texas, USA, volume 1730 of Lecture Notes in Artificial
Intelligence, pages 357–371. Springer, Berlin, 1999.

[HS0x]   Pascal Hitzler and Anthony K. Seda. Generalized metrics and uniquely determined
           logic programs. Theoretical Computer Science, 200x. To appear.

[Lif99]   Vladimir Lifschitz. Answer set planning. In Danny De Schreye, editor, Logic Pro-
           gramming. Proceedings of the 1999 International Conference on Logic Program-
           ming, pages 23–37, Cambridge, Massachusetts, 1999. MIT Press.

[Llo88]   John W. Lloyd. Foundations of Logic Programming. Springer, Berlin, 1988.

[LMR92]   Jorge Lobo, Jack Minker, and Arcot Rajasekar. Foundations of disjunctive logic
           programming. MIT Press, 1992.

[LRS97]   Nicola Leone, Pasquale Rullo, and Francesco Scarello. Disjunctive stable models:
           Unfounded sets, fixpoint semantics, and computation. Information and Computa-
           tion, 135(2):69–112, 1997.

[Mar96]   Elena Marchiori. On termination of general logic programs with respect to con-
           structive negation. Journal of Logic Programming, 26(1):69–89, 1996.

[MT99]   V. Wiktor Marek and Miroslav Truszczyński. Stable models and an alternative
           logic programming paradigm. In Krzysztof R. Apt, V. Wiktor Marek, Miroslav
           Truszczyński, and David S. Warren, editors, The Logic Programming Paradigm: A
           25 Year Persepective, pages 375–398. Springer, Berlin, 1999.

[PP90]   Halina Przymusinska and Teodor C. Przymusinski. Weakly stratified logic pro-
           grams. Fundamenta Informaticae, 13:51–65, 1990. Krzysztof R. Apt, editor, special
           issue of Fundamenta Informaticae on Logical Foundations of Artificial Intelligence.

[Prz88]   Teodor C. Przymusinski. On the declarative semantics of deductive databases and
           logic programs. In Jack Minker, editor, Foundations of Deductive Databases and
           Logic Programming, pages 193–216. Morgan Kaufmann, Los Altos, CA, 1988.

[RZ01]   William C. Rounds and Guo-Qiang Zhang. Clausal logic and logic programming
           in algebraic domains. Information and Computation, 171(2):156–182, 2001.

[Sed95]   Anthony Karel Seda. Topology and the semantics of logic programs. Fundamenta
           Informaticae, 24(4):359–386, 1995.

[SHLG94]   Viggo Stoltenberg-Hansen, Ingrid Lindström, and Edward R. Griffor. Mathemat-
           ical Theory of Domains. Cambridge University Press, 1994.

[SNS0x]   Patrik Simons, Ilkka Niemelä, and Timo Soininen. Extending and implementing
           the stable model semantics. Artificial Intelligence, 200x. To appear.

[Sub99]   V.S. Subrahmanian. Nonmonotonic logic programming. IEEE Transactions on
           Knowledge and Data Engineering, 11(1):143–152, January/February 1999.

[VGRS91] Allen Van Gelder, Kenneth A. Ross, and John S. Schlipf. The well-founded se-
mantics for general logic programs. *Journal of the ACM*, 38(3):620–650, 1991.

[Wen02] Matthias Wendt. Towards a unified view of the hierarchy of logic program classes.
Project Thesis, Knowledge Representation and Reasoning Group, Artificial Intel-
ligence Institute, Department of Computer Science, Dresden University of Tech-
nology, 2002. To appear.

[WZ00] Haixun Wang and Carlo Zaniolo. Nonmonotonic reasoning in LDL^{++}. In Jack
Minker, editor, *Logic-Based Artificial Intelligence*, chapter 22, pages 523–542.
Kluwer Academic Publishers, Boston, 2000.