Adding Closed World Assumptions to Well Founded Semantics

Luís Moniz Pereira José J. Alferes
Joaquim N. Aparício
AI Centre, Uninova and DCS, U. Nova de Lisboa
2825 Monte da Caparica, Portugal
{lmp,jja,jna}@fct.unl.pt

Abstract

Given a program P we specify an enlargement of its Well Founded Model which gives meaning to the adding of Closed World Assumptions. We do so by proposing the desirable principles of a Closed World Assumption (CWA), and proceed to formally define and apply them to Well Founded Semantics (WFS), in order to obtain a WFS added with CWA, the O-semantics. After an introduction and motivating examples, there follow the presentation of the concepts required to formalize the model structure, the properties it enjoys, and the criteria and procedures which allow the precise characterization of the preferred unique maximal model that gives the intended meaning to the O-Semantics of a program, the O-Model. Some properties are also exhibited that permit a more expedite obtention of the models. Several detailed examples are introduced throughout to illustrate the concepts and their application. Comparison is made with other work, and in the conclusions the novelty of the approach is brought out.

1 Introduction and Motivation

Well Founded Semantics [14] has been proposed as a suitable semantics for general logic programs. Its Extended Stable Models (XSM) [12, 13] version, and the inclusion of a second type of negation, have been explored as a framework for formalizing a variety of forms of non-monotonic reasoning [10, 11] and generalized to deal with contradiction removal and counterfactuals [7, 8, 9]. The increasing rôle of logic programming extensions as an encompassing framework for these and other AI topics is expounded at

length in [4], where they argue, and we concur, that WFS is by design overly careful in deciding about the falsity of some atoms, leaving them undefined, and that a suitable form of CWA can be used to safely and undisputably assume false some of the atoms absent from the well founded model of a program. Consider the following example adapted from [3], itself a variant of the "game" example of [2]:

Example 1 Consider the program:

$$\begin{aligned} win(X) &\leftarrow move(X,Y), {\sim}win(Y) \\ raisedBet(X) &\leftarrow win(X) \\ move(a,a) &\leftarrow \\ move(b,c) &\leftarrow \end{aligned}$$

expressing that "X is a winning position if there is a move from X to Y and Y is not a winning position", "in a winning position bets are raised", and that "we can move from position a to position a, and from position b to c".

c is not a winning position since it is impossible to move from c. b is a winning position because it is possible to move from b to c and c is not a winning position. a is a position of draw.

Neither $win(a)$ nor ${\sim}win(a)$ should hold. This is correctly handled by WFS which assigns the truth–value undefined to $win(a)$.

The semantics of this program should also capture the intended meaning that bets are not raised in a position of draw. This is not captured by WFS which leave $raisedBet(a)$ undefined. □

More abstractly, let $P = \{a \leftarrow {\sim}a;\ c \leftarrow a\}$, where $WFM(P) = \{\}$. We argue that the intended meaning of the program may be $\{{\sim}c\}$, since a may not be true in any extended stable model of P, by the first rule, and so, the second rule cannot contradict the assigned meaning. Another way to understand this is that one may safely assume ${\sim}c$ using a form of CWA on c, since ${\sim}a$ may not be consistently assumed.

However, when relying on the absence of present evidence about some atom A, we do not always want to assume that ${\sim}A$ holds, since there may exist consistent assumptions allowing to conclude A. Roughly, we want to define the notion of concluding for the truth of a negative literal ${\sim}A$ just in case there is no hard nor hypothetical evidence to the contrary, i.e. no consistent set of negative assumptions such that ${\sim}A$ is untenable.

Consider $P = \{a \leftarrow {\sim}b;\ b \leftarrow {\sim}a;\ c \leftarrow a\}$. If we interpret the meaning of this program as its WFM (which is empty), and as we do not have a, a naïve

CWA could be tempted to derive $\sim c$ based on the assumption $\sim a$. There is however an alternative negative assumption $\sim b$ that, if made, defeats the assumption $\sim a$, i.e. the assumption $\sim a$ may not be sustained since it can be defeated by the assumption $\sim b$. We will define later more precisely the notions of sustainability and tenability.

Both programs above have empty well founded models. We argue that WFS is too careful, and something more can safely be added to the meaning of program, thus reducing the undefinedness of the program, if we are willing to adopt a suitable form of CWA.

We contend that a set $CWA(P)$ of negative literals (assumptions) *added* to a program model $MOD(P)$ by CWA must obey the four principles:

1. $MOD(P) \cup CWA(P) \not\models L$ for any $\sim L \in CWA(P)$. This says that the program model added with the set of assumptions identified by the CWA rule must be **consistent**.

2. There is no other set of assumptions A such that $MOD(P) \cup A \models L$ for some $\sim L \in CWA(P)$. I.e. $CWA(P)$ is **sustainable**.

3. $CWA(P)$ must be **unique**.

4. $CWA(P)$ must, additionally, be **maximal**.

The paper is organized as follows: in the next section we present some basic definitions. In section 3 we introduce some new definitions, capturing the concepts behind the semantics, accompanied by examples illustrating them. Models are defined and organized into a lattice, and the class of sustainable A-Models is identified. In section 5 we define the O-Semantics of a program P based on the class of maximal sustainable tenable A-Models. A unique model is singled out as the O-Model of P. Afterwards we present some properties of the class of A-Models. Finally, we relate to other semantics and present conclusions.

2 Language

Here we give basic definitions and establish notation ([6]). A program is a set of rules of the form:

$$H \leftarrow B_1, \ldots, B_n, \sim C_1, \ldots, \sim C_m \ (n \geq 0, m \geq 0)$$

or equivalently

$$H \leftarrow \{B_1, \ldots, B_n\} \cup \sim\{C_1, \ldots, C_m\},$$

where $\sim\{A_1, \ldots, A_n\}$ is a shorthand for $\{\sim A_1, \ldots, \sim A_n\}$, and $\sim C$ is short for $\sim\{C1, \ldots, C_m\}$; H, B_i and C_j are atoms.

The Herbrand Base $B(P)$ of a program P is defined as usual as the set of all ground atoms. An interpretation I of P is denoted by $T \cup \sim F$, where T and F are disjoint subsets of $B(P)$. Atoms in T are said to be *true* in I, atoms in F *false* in I, and atoms in $B(P) - (T \cup F)$ *undefined* in I.

In an interpretation $T \cup \sim F$ a conjunction of literals

$$\{B_1, \ldots, B_n\} \cup \sim\{C_1, \ldots, C_m\}$$

is true	iff	$\{B_1, \ldots, B_n\} \subseteq T$ and $\{C_1, \ldots, C_m\} \subseteq F$
is false	iff	$\{B_1, \ldots, B_n\} \cap F \neq \emptyset$ or $\{C_1, \ldots, C_m\} \cap T \neq \emptyset$
is undefined	iff	it is neither true nor false.

3 Adding Negative Assumptions to a Program

Here we show how to consistently add negative assumptions to a program P. Informally, it is consistent to add a negative assumption to P if the assumption atom is not among the consequences P after adding the assumption. We also define when a set of negative assumptions is defeated by another, and show how the models of a program, for different sets of negative assumptions added to it, are organized into a lattice.

We begin by defining what it means to add assumptions to a program. This is achieved by substituting *true* for the assumptions, and *false* for their atoms, in the body of all rules.

Definition 3.1 (P+A) The program $P + A$ obtained by adding to a program P a set of negative assumptions $A \subseteq \sim B(P)$ is the result of:

- Deleting all rules $H \leftarrow \{B_1, \ldots, B_n\} \cup \sim C$ from P such that some $\sim B_i \in A$
- Deleting from the remaining rules all $\sim L \in A$

Definition 3.2 (Assumption Model) An Assumption Model of a program P, or A-Model for short, is a pair $\langle A; M \rangle$ where $A \subseteq \sim B(P)$ and $M = WFM(P + A)$.

Among these models we define the partial order $\leq_a$ in the following way: $\langle A_1; M_1\rangle \leq_a \langle A_2; M_2\rangle$ iff $A_1 \subseteq A_2$. On the basis of set union and set intersection among the sets A of negative assumptions, the set of all A-Models becomes organized as a complete lattice.

Having defined assumption models we next consider their consistency. According to the CWA principles above, an assumption $\sim A$ cannot be added to a program P if by doing so A is itself a consequence of P, or some other assumption is contradicted.

Definition 3.3 (Consistent A-Model) An A-Model $\langle A; M\rangle$ is consistent iff $A \cup M$ is an interpretation, i.e. there exists no assumption $\sim L \in A$ such that $L \in M$.

Example 2 Let P :

$$\begin{array}{lcl} c & \leftarrow & \sim b \\ b & \leftarrow & \sim a \\ a & \leftarrow & \sim a \end{array}$$

whose WFM is empty. The A-Model $\langle\{\sim a\}; \{a, b, \sim c\}\rangle$ is inconsistent since by adding the assumption $\sim a$ then $a \in WFM(P + \{\sim a\})$. The same happens with all A-Models containing the assumption $\sim a$. The A-Model $\langle\{\sim b, \sim c\}; \{c\}\rangle$ is also inconsistent. Thus the only consistent A-Models are $\langle\{\}; \{\}\rangle$, $\langle\{\sim b\}; \{c\}\rangle$ and $\langle\{\sim c\}; \{\}\rangle$. □

Lemma 3.1 *If an A-Model AM is inconsistent then any A-Model AM′ such that $AM \leq_a AM'$ is inconsistent.*

Proof: We prove that for all $\sim a' \in B(P)$, if $\langle A; WFM(P + A)\rangle$ is inconsistent then $\langle A \cup \{\sim a'\}; WFM(P + (A \cup \{\sim a'\}))\rangle$ is also inconsistent. By definition of inconsistent A-Model: $\exists \sim b \in A \mid b \in WFM(P + A)$, so it suffices to guarantee that :

$$b \notin WFM(P + (A \cup \{\sim a'\})) \Rightarrow a' \in WFM(P + (A \cup \{\sim a'\})).$$

Consider $b \notin WFM(P + A \cup \{\sim a'\})$. Since b is true in $P + A$, and since $P + (A \cup \{\sim a'\})$ only differs from $P + A$ in rules with a' or $\sim a'$ in the body, it follows that there is a support set $SS_{P+A}(b)$ containing a' (in appendix A we recall the definition of support set introduced in [7]), and thus, by definition of support set, a' is also true in $P + A$.

Since $a' \in WFM(P+A)$, by propositions A.1 and A.2 there is a support set $SS_{P+A}(a')$ such that $a' \notin SS_{P+A}(a')$ and $\sim a' \notin SS_{P+A}(a')$. As the addition of $\sim a'$ to $P + A$ only changes rules with $\sim a'$ or a', then:

$$Rules(SS_{P+A}(a')) \subseteq P + (A \cup \{\sim a'\}) \subseteq P + A$$

and by proposition A.3 $a' \in WFM(P + (A \cup \{\sim a'\}))$. $\diamondsuit$

According to the CWA principles above, an assumption $\sim A$ cannot be sustained if there is some set of consistent assumptions that concludes A. We've already expressed the notion of consistency being used. To capture the notion of sustainability we now formally define how an A-Model can defeat another, and define sustainable A-Models as the nondefeated consistent ones.

Definition 3.4 (Defeating) *A consistent A-Model $\langle A; M\rangle$ is defeated by a consistent $\langle A'; M'\rangle$ iff $\exists \sim a \in A | a \in M'$.*

Definition 3.5 (Sustainable A-Models) *An A-Model $\langle A; M\rangle$ is sustainable iff it is consistent and not defeated by any consistent A-Model. Equivalently $\langle \sim S; M\rangle$ is sustainable iff:*

$$S \cap \bigcup_{consistent\ \langle A_i; M_i\rangle} M_i = \{\}$$

Example 3 The only sustainable models in example 2 are $\langle\{\}; \{\}\rangle$ and $\langle\{\sim b\}; \{c\}\rangle$. Note that the consistent A-Model $\langle\{\sim c\}; \{\}\rangle$ is defeated by $\langle\{\sim b\}; \{c\}\rangle$, i.e. the assumption $\sim c$ is unsustainable since there is a set of consistent assumptions (namely $\{\sim b\}$) that leads to the conclusion c. $\square$

The assumptions part of maximal sustainable A-Models of a program P are maximal sets of consistent Closed World Assumptions that can be safely added to the consequences of P without risking contradiction by other assumptions.

Lemma 3.2 *If an A-Model AM is defeated by another A-Model D, then all A-Models AM' such that $AM \leq_a AM'$ are defeated by D.*

Proof: If $AM = \langle A; M\rangle$ is defeated by $D = \langle A_D; M_D\rangle$, then there exists $d \in M_D$ such that $\sim d \in A$. Since all AM's are of the form $AM' = \langle A'; M'\rangle$ where $A' = A \cup B$ then $\sim d \in A'$, i.e. D defeats AM'. $\diamondsuit$

Lemma 3.3 *The A-Model* $\langle\{\}; WFM(P)\rangle$ *is always sustainable.*

Proof: By definition of sustainable. $\diamondsuit$

Theorem 3.4 *The set of all sustainable A-Models of a program is nonempty. On the basis of set union and set intersection among their A sets, the A-Models ordered by* $\leq_a$ *form a lower semi lattice.*

Proof: Follows directly from the above lemmas. $\diamondsuit$

A program may have several maximal sustainable A-Models.

Example 4 Let P be:

$$\begin{array}{lcl} c & \leftarrow & \sim c, \sim b \\ b & \leftarrow & a \\ a & \leftarrow & \sim a \end{array}$$

Its sustainable A-Models are $\langle\{\}; \{\}\rangle$, $\langle\{\sim b\}; \{\}\rangle$ and $\langle\{\sim c\}; \{\}\rangle$. The last two are maximal sustainable A-Models. We cannot add both $\sim b$ and $\sim c$ to the program to obtain a sustainable A-Model since $\langle\{\sim b, \sim c\}; \{c\}\rangle$ is inconsistent. $\square$

4 The O-semantics

This section is concerned with the problem of singling out, among all sustainable A-Models of a program P, one that uniquely determines the meaning of P when the CWA is enforced. This is accomplished by means of a selection criterium that takes a lower semilattice of sustainable A-Models and obtains a subsemilattice of it, by deleting A-Models that in a well–defined sense are less preferable, i.e. the untenable ones.

Sustainability of a consistent set of negative assumptions insists that there be no other consistent set that defeats it (i.e. there is no hypothetical evidence whose consequences contradict the sustained assumptions). Tenability requires that a maximal sustainable set of assumptions be not contradicted by the consequences of adding to it another competing (nondefeating and nondefeated) maximal sustainable set.

The selection process is repeated and ends up with a complete lattice of sustainable A-Models, which defines for every program P its O-Semantics. The meaning of P is then specified by the greatest A-Model of the semantics, its O-Model.

To illustrate the problem of preference among maximal A-Models we introduce an example.

Example 5 Consider the program P :

$$\begin{array}{lcl} c & \leftarrow & \sim c, \sim b \\ b & \leftarrow & a \\ a & \leftarrow & \sim a \end{array}$$

whose sustainable A-Models are $\langle\{\};\{\}\rangle$, $\langle\{\sim b\};\{\}\rangle$, and $\langle\{\sim c\},\{\}\rangle$. Because we wish to maximize the number of negative assumptions, we consider the maximal A-Models, which in this case are the last two. The join of these maximal A-Models, $\langle\{\sim b, \sim c\};\{c\}\rangle$, is per force inconsistent, in this case wrt c. This means that when assuming $\sim c$ there is an additional set of assumptions entailing c, making this A-Model *untenable*. But the same does not apply to $\sim b$. Thus the preferred A-Model is $\langle\{\sim b\},\{\}\rangle$, and the A-Model $\langle\{\sim c\};\{\}\rangle$ is said untenable. The rationale for the preference is grounded on the fact that the inconsistency of the join arises wrt c but not wrt b. □

Definition 4.1 (Candidate Structure) *A Candidate Structure CS of a program P is any subsemilattice of the lower semi lattice of all sustainable A-Models of P.*

Definition 4.2 (Untenable A-Models) *Let* $\{\langle A_1; M_1\rangle, \ldots, \langle A_n; M_n\rangle\}$ *be the set of all maximal A-Models in Candidate Structure CS. Let* $J = \langle A_J; M_J\rangle$ *be the join of all such A-Models, in the complete lattice of all A-Models. An A-Model* $\langle A_i; M_i\rangle$ *is untenable wrt CS iff it is maximal in CS and there exists* $\sim a \in A_i$ *such that* $a \in M_J$.

Proposition 4.1 *There exists no untenable A-Model wrt a Candidate Structure with a single maximal element.*

Proof: Since the join coincides with the unique maximal A-Model, which is sustainable by definition of CS, then it cannot be untenable. ♢

The Candidate Structure left after removing all untenable A-Models of a CS may itself have several untenable elements, some of which might not be untenable A-Models in the initial CS. If the removal of untenable A-Models is performed repeatedly on the retained Candidate Structure, a structure with no untenable models is eventually obtained, albeit the bottom element of the CSs.

Definition 4.3 (Retained CS) *The Retained Candidate Structure $R(CS)$ of a Candidate Structure CS is defined recursively in the following way (where J is the join of elements of CS in the complete lattice of all A-Models):*

- *$J \cup CS$ if there are no untenable A-Models in CS*
- *Otherwise, let Unt be the set of all untenable A-Models wrt CS. Then $R(CS) = R(CS - Unt)$*

Definition 4.4 (The O-Semantics) *The O-Semantics of a program P is defined by the Retained Candidate Structure of the semilattice of all sustainable A-Models of P.*

Let $\langle A; M\rangle$ be its maximal element. The intended meaning of P is $A \cup M$, the O-Model of P.

Remark 4.1 At this point, we are in a position to make an important remark. Our goal is to maximally reduce undefinedness of the well–founded model by adding to it negative assumptions. Now, the peeling process of subtracting only *maximal* untenable A-Models from candidate structures ends up with a retained candidate structure with a maximal element. So we must guarantee this element is always greater or equal than the result we would obtain if we didn't require untenable A–Models to be *maximal* in definition 4.2.

This is indeed guaranteed, for the join of the maximal elements of each candidate structure is always greater than any join of non–maximal elements of that structure, and because the maximal element of the retained lattice is by definition one such join of maximal elements.

Example 10 shows that if untenable A–Models were not defined as maximal then a smaller O–Model would be obtained.

Theorem 4.1 *The O-semantics of a program P is always defined by a complete lattice of sustainable A-Models.*

Proof: Since every candidate structure is a semilattice of sustainable A-Models, it is enough to prove that the join $J = \langle A_J; M_J\rangle$ of the retained candidate structure CS of the semilattice of all sustainable A-Models of P is a sustainable A-Model.

If we assume that J is inconsistent then at least one maximal A-Model in CS is untenable. Accordingly, since in the final retained CS there are, by definition, no untenable A-Models, J is consistent.

J cannot be defeated by any other consistent A-Model D because, in such a case, at least one other element of CS would also be deafeated by D, which is impossible by definition of candidate structure. $\diamondsuit$

Corollary 1 *The O-Semantics of a program has no untenable A-Model wrt itself.*

Proof: Follows directly from the theorem and proposition 4.1. $\diamondsuit$

Corollary 2 (Existence of the O-Semantics) *The Retained Candidate Structure of the semilattice of all sustainable A-Models is nonempty.*

Proof: Follows directly from the theorem. $\diamondsuit$

5 Examples

In this section we display some examples and their O-Semantics. Remark that indeed the O-Models obtained express the safe CWAs compatible with the WFMs (which are all {}). In section 7, "Relation to other work" additional examples can be found which bring out the distinctness of O–semantics wrt other semantics.

Example 6 Let P be:

$$\begin{array}{lcl} a & \leftarrow & \sim a \\ b & \leftarrow & a \\ c & \leftarrow & \sim c, \sim b \\ d & \leftarrow & c \end{array}$$

The semilattice of all sustainable A-Models CS is shown in figure 1.

The join of its maximal A-Models is $\langle\{\sim b, \sim c, \sim d\}; \{c, \sim d\}\rangle$. Consequently, the maximal A-Model on the right is untenable since it contains $\sim c$ in the assumptions, and c is a consequence of the join. So $R(CS) = R(CS')$ where CS' is as in figure 2.

The join of all maximal elements in CS' is the same as before and the only untenable A-Model is again the maximal one having $\sim c$ in its assumptions. Thus $R(CS) = R(CS'')$ where CS'' is shown in figure 3.

So the O-Model is $\{\sim b, \sim d\}$. Note that if P is divided into $P_1 = \{c \leftarrow \sim c, \sim b;\ d \leftarrow c\}$ and $P_2 = \{a \leftarrow \sim a;\ b \leftarrow a\}$, the O-models of P_1 and P_2 both agree on the only common literal $\sim b$. So $\sim b$ rightly belongs to the O-models of P. $\square$

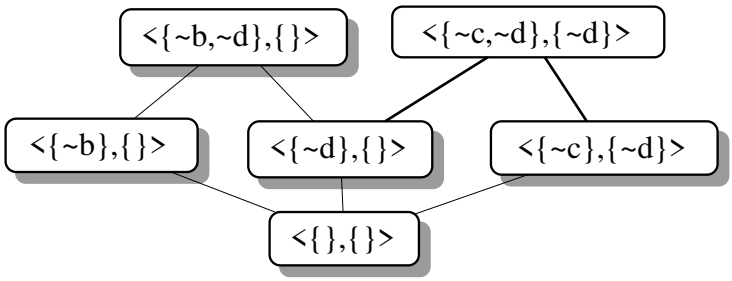

Figure 1: Semilattice of Sustainable A-Models of example 6. Tenable A-Models are shadowed.

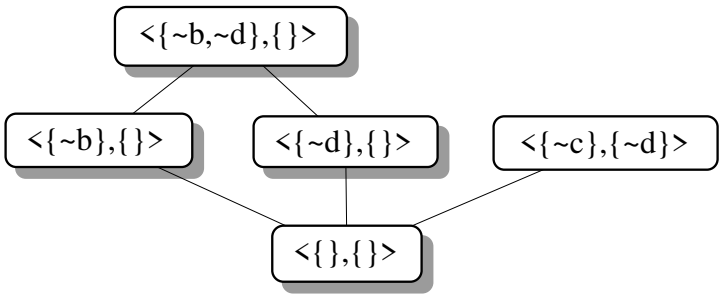

Figure 2: CS' as explained in example 6.

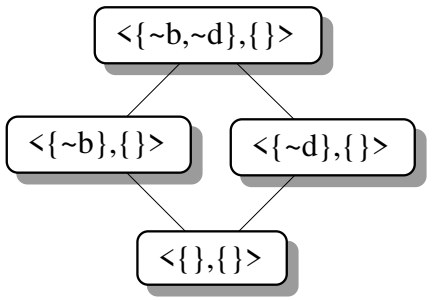

Figure 3: Retained candidate structure of example 6.

Example 7 Consider the P :

$$\begin{array}{lcl} q & \leftarrow & \sim p \\ p & \leftarrow & a \\ a & \leftarrow & \sim b \\ b & \leftarrow & \sim c \\ c & \leftarrow & \sim a \end{array}$$

Its only consistent A-Models are $\langle\{\};\{\}\rangle$, $\langle\{\sim p\};\{q\}\rangle$ and $\langle\{\sim q\};\{\}\rangle$. As this last one is defeated by the second, the only sustainable ones are the first two. Since only one is maximal, these two A-Models determine the O-Semantics, and the meaning of P is $\{\sim p, q\}$, its O-Model. Note that if the three last rules, forming an *"undefined loop"*, are replaced by another *"undefined loop"* $a \leftarrow\sim a$, the O-model is the same. This is as it should, since the first two rules conclude nothing about a. □

Example 8 Let P :

$$\begin{array}{lcl} p & \leftarrow & a,\ b \\ a & \leftarrow & \sim b \\ b & \leftarrow & \sim a \end{array}$$

The A-Models with $\sim b$ in their assumptions defeat A-Models with $\sim a$ in their assumptions and vice-versa. Thus the O-Semantics is determined by $\langle\{\};\{\}\rangle$ and $\langle\{\sim p\};\{\}\rangle$, and the meaning of P is $\{\sim p\}$, its O-Model. □

Example 9 Consider the program P :

$$\begin{array}{lcl} c & \leftarrow & \sim c,\ \sim b \\ b & \leftarrow & \sim c,\ \sim b \\ b & \leftarrow & a \\ a & \leftarrow & \sim a \end{array}$$

Its sustainable A-Models are $\langle\{\};\{\}\rangle$, $\langle\{\sim b\};\{\}\rangle$ and $\langle\{\sim c\};\{\}\rangle$. The join of the two maximal ones is $\langle\{\sim b, \sim c\};\{b, c\}\rangle$, and so both are untenable. Thus the Retained Candidate Structure has the single element $\langle\{\};\{\}\rangle$ and the meaning of P is $\{\}$. □

Example 10 Consider the program P :

$$\begin{array}{lll} c & \leftarrow & \sim a,\ \sim c \\ c & \leftarrow & \sim b,\ \sim c \\ a & \leftarrow & \sim b,\ \sim c \\ a & \leftarrow & d \\ b & \leftarrow & d \\ c & \leftarrow & d \\ d & \leftarrow & \sim d \end{array}$$

The semilattice of all sustainable A-Models CS is as in figure 4.

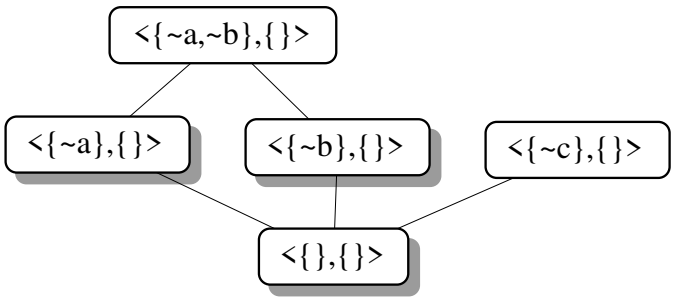

Figure 4: Semilattice of sustainable A-Models of example 10.

The join of its maximal A-Models is $\langle\{\sim a, \sim b, \sim c\}; \{a, c\}\rangle$. Consequently, all maximal A-Models are untenable. So $R(CS) = R(CS')$ where CS' is shown in figure 5.

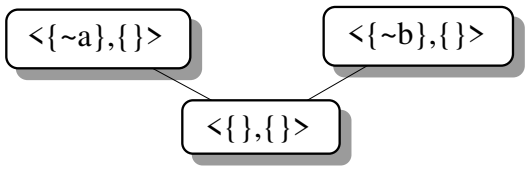

Figure 5: CS' as explained in example 10.

Since the join of all elements is $\langle\{\sim a, \sim b\}; \{\}\rangle$, there are no untenables in CS'. Thus the O-Semantics is as depicted in figure 6.

and the O-Model is $\{\sim a, \sim b\}$.

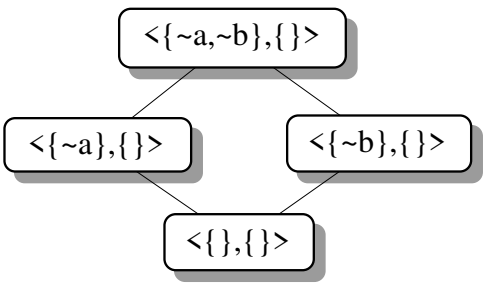

Figure 6: The O–Semantics of example 10.

If untenable A-Models were not defined as maximal ones (cf. remark 4.1) then $\langle\{\sim a\};\{\}\rangle$ would also be untenable wrt to the semilattice of all sustainable A-Models CS. Then the $R(CS)$ would be as presented in figure 7, and the O-Model would be smaller: $\{\sim b\}$.. □

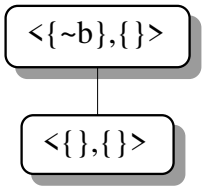

Figure 7: Retained CS of example 10, if untenable A-Models were not defined as maximal ones.

6 Properties of Sustainable A-Models

This section explores properties of sustainable A-Models that provide a better understanding of them, and also give hints for their construction without having to previously calculate all A-Models.

We begin with properties that show how our models can be viewed as an extension to Well Founded Semantics (WFS). As mentioned in [6], negation in WFS is based on the notion of support, i.e. a literal $\sim L$ only belongs to an Extended Stable Model (XSM) if all the rules for L (if any) have false bodies in the XSM. In contradistinction, we are interested in negations

as consistent hypotheses that cannot be defeated. To that end we weaken the necessary (but not sufficient) conditions for a negative literal to belong to a model as explained below. We still want to keep the necessary and sufficient conditions of support for positive literals. More precisely, knowing that XSMs must obey, among others, the following conditions cf. [6]:

- If there exists a rule $p \leftarrow B$ in the program such that B is true in model M then p is also true in M (*sufficiency of support for positive literals*).
- If an atom $p \in M$ then there exists a rule $p \leftarrow B$ in the program such that B is true in M (*necessity of support for positive literals*).
- If all rule bodies for p are false in M then $\sim p \in M$ (*sufficiency of support for negative literals*).
- If $\sim p \in M$ then all rules for p have false bodies in M (*necessity of support for negative literals*).

Our consistent A-models, when understood as the union of their pair of elements, assumptions A and $WFM(P + A)$, need not obey the fourth condition. Foregoing it condones making negative assumptions. In our models an atom might be false even if it has a rule whose body is undefined. Thus, only false atoms with an undefined rule body are candidates for having their negation added to the $WFM(P)$.

Proposition 6.1 *Let $\langle A; M\rangle$ be any consistent A-Model of a program P. The interpretation $A \cup M$ obeys the first three conditions above.*

Proof: Here we prove the satisfaction of the first condition. The remaining proofs are along the same lines. If:

$$\exists p \leftarrow b_1, \ldots, b_n, \sim c_1, \ldots, \sim c_m \in P \mid \{b_1, \ldots, b_n, \sim c_1, \ldots, \sim c_m\} \subseteq A \cup M$$

then

$$b_i \in M \ (1 \leq i \leq n) \ and \ \sim c_j \in M \ or \ \sim c_j \in A \ (1 \leq j \leq m).$$

Let $p \leftarrow b_1, \ldots, b_n, \sim c_l, \ldots, \sim c_k (l \geq 1, k \leq m)$ be the rule obtained from an existing one by removing all $\sim c_j \in A$, which is, by definition, a rule of $P+A$. Thus there exists a rule $p \leftarrow B$ in $P+A$ such that $B \subseteq WFM(P+A) = M$.

Given that the WFM of any program must obey the first condition above, $p \in WFM(P + A)$. $\diamondsuit$

Next we state properties useful for more directly finding the sustainable A-Models.

Proposition 6.2 *There exists no consistent A-Model $\langle A; M\rangle$ of P with $\{\sim a\} \subseteq A$ such that $a \in WFM(P)$.*

Proof: We begin by proving the proposition for $\{\sim a\} = A$.

Since $a \in WFM(P)$, then by propositions A.1 and A.2 there is a $SS_P(a) = S$ such that $a \notin S$ and $\sim a \notin S$, and consequently

$$Rules(S) \subseteq P + \{\sim a\}.$$

Then, by proposition A.3, $a \in WFM(P + \{\sim a\})$, and thus $\langle \{\sim a\}; WFM(P + \{\sim a\})\rangle$ is inconsistent.

It follows, from lemma 3.1, that all A-Models $\langle A; M\rangle$ such that $\{\sim a\} \subseteq A$ are inconsistent. $\diamondsuit$

Hence, A-Models not obeying the above restriction are not worth considering as sustainable.

Proposition 6.3 *If a negative literal $\sim L \in WFM(P)$ then there is no consistent A-Model $\langle A; M\rangle$ of P such that $L \in M$.*

Proof: We prove that if $L \in M$ for a given A-Model $\langle A; M\rangle$ of P then $\langle A; M\rangle$ is inconsistent. If $L \in M$ there must exist a rule $L \leftarrow B, \sim C$ in P such that $B \cup \sim C \subseteq M \cup A$ and $B \cup \sim C$ is false in $WFM(P)$, i.e. there must exist $L \leftarrow B, \sim C$ in P with at least one body literal true in $M \cup A$ and false in $WFM(P)$. If that literal is an element of $\sim C$, by proposition 6.2 $\langle A; M\rangle$ is inconsistent (its corresponding atom is true in $WFM(P)$ and false in $M \cup A$). If it is an element of B this theorem applies recursively, ending up in a rule with empty body, an atom with no rules or a loop without an interposing $\sim l$. As shown below the truth value of literals in these conditions can never be changed:

Since the $P + A$ operation only involves deleting rules with literals at the body and literals from the body of rules, the truth value of atoms without rules is always false no matter which A is being considered, and the truth value of atoms with a fact is always false. Literals in a loop without interposing $\sim l$ are false in P, and remain false if rules of the loop are deleted.

$\diamondsuit$

Theorem 6.1 *If* $\sim L \in WFM(P)$ *then* $\sim L \in M$ *in every consistent A-Model* $\langle A; M \rangle$ *of* P.

Proof: Given proposition 6.3, it suffices to prove that L is not undefined in any consistent A-Model of P. The proof is along the lines of that of the proposition above. $\diamondsuit$

Consequently, all negative literals in the $WFM(P)$ belong to every sustainable A-Model.

Lemma 6.2 *Let* $WFM(P) = T \cup \sim F$. *For any subset* S *of* $\sim F$

$$WFM(P) = WFM(P + S).$$

Proof: This lemma is easily shown using the definition of $P + A$ and the properties of the WFM. $\diamondsuit$

Theorem 6.3 *Let* $WFM(P) = T \cup \sim F$ *and* $\langle A; WFM(P + A) \rangle$ *be a consistent A-Model, and let* $A' = A \cap \sim F$. *Then*

$$WFM(P + A) = WFM(P + (A - A')).$$

Proof: Let $P' = P + (A - A')$ and $WFM(P) = T \cup \sim F$.
By theorem 6.1 $\sim F \subseteq WFM(P')$. So, by lemma 6.2,

$$WFM(P') = WFM(P' + \sim F) = WFM([P + (A - (A \cap \sim F))] + \sim F).$$

By definition of $P + A$, it follows that $(P + A_1) + A_2 = P + (A_1 \cup A_2)$. Thus $WFM(P')$ is:

$$WFM(P + [(A - (A \cap \sim F)) \cup \sim F]) = WFM(P + A)$$

$\diamondsuit$

N.B. This theorem shows that sets of assumptions including negative literals of $WFM(P)$ are not worth considering since there exist smaller sets having exactly the same consequences $A \cup M$ and, by proposition 6.3, the larger sets are not defeatable by reason of negative literals from the $WFM(P)$. Hence, in the remainder of the paper, we consider only A–Models

whose assumptions are not in the WFM, inasmuch all $WFM(P)$ assumptions will be part of $WFM(P + A)$ for any A.

Another important hint for calculating the sustainable A-Models is given by lemma 3.1. According to it one should start by calculating A-Models with smaller assumption sets, so that when an inconsistent A-Model is found, by the lemma, sets of assumptions containing it are unworthy of consideration.

Example 11 Let P :

$$\begin{array}{lcl} p & \leftarrow & \sim a, \sim b \\ a & \leftarrow & c, d \\ c & \leftarrow & \sim c \\ d & & \end{array}$$

The least A-Model is $\langle\{\}; \{d, \sim b\}\rangle$ where $\{d, \sim b\} = WFM(P)$. Thus sets of assumptions containing $\sim d$ or $\sim b$ are not worth considering. Take now, for example, the consistent A-Model $\langle\{\sim a\}; \{d, \sim b, p\}\rangle$, which we retain. Consider $\langle\{\sim c\}; \{c, a, \sim p\}\rangle$; as this A-Model is inconsistent we do not retain it nor consider any other A-Models with assumption sets containing $\sim c$. Now we are left with just two more A-Models worth considering: $\langle\{\sim p\}; \{d, \sim b\}\rangle$ which is defeated by $\langle\{\sim a\}; \{d, \sim b, p\}\rangle$; and $\langle\{\sim p, \sim a\}; \{d, \sim b, p\}\rangle$ which is inconsistent. Thus the only two sustainable A-Models are $\langle\{\}; \{d, \sim b\}\rangle$ and $\langle\{\sim a\}; \{d, \sim b, p\}\rangle$. In this case, the latter is the single maximal sustainable A-Model, and thus uniquely determines the intended meaning of P to be $A \cup M = \{\sim a, d, \sim b, p\}$. □

7 Relation to other work

Consider the following program P :

$$\begin{array}{lcl} p & \leftarrow & q, \sim r, \sim s \\ q & \leftarrow & r, \sim p \\ r & \leftarrow & p, \sim q \\ s & \leftarrow & \sim p, \sim q, \sim r \end{array}$$

In [12] they argue that the intended semantics of this program should be the interpretation $\{s, \sim p, \sim q, \sim r\}$, due to the mutual circularity of p, q, r. This model is precisely the meaning assigned to the program by the O-Semantics, its O-Model. Note that WFS identifies the (3-valued) empty model as the meaning of the program. This is also the model provided by

stable model semantics [2]. The weakly perfect model semantics for this program is undefined as noticed in [12].

The EWFS [1] is also an extension to the WFM based on the notion of GCWA [5]. Roughly, EWFS moves closer than the WFM (in the sense of being less undefined) to being the intersection of all minimal Herbrand models of P. With a different notation from that of [1]:

$$EWFM(P) =_{def} WFM(P) \cup T(WFM(P)) \cup \sim F(WFM(P))$$

where:

$$T(\mathcal{I}) =_{def} True(MIN_MOD(\mathcal{I}, P)),$$
$$F(\mathcal{I}) =_{def} False(MIN_MOD(\mathcal{I}, P))$$

$\mathcal{I}$ is a three–valued interpretation, and $MIN_MOD(\mathcal{I}, P)$ is the collection of all minimal two–valued Herbrand models of P consistent with $\mathcal{I}$. For a set $\mathcal{S}$ of interpretations, $True(\mathcal{S})$ (resp. $False(\mathcal{S})$) denotes the set of all atoms which are true (resp. false) in all interpretations of $\mathcal{S}$.

For the program $P = \{a \leftarrow\sim a\}$ we have:

$$WFM(P) = \{\},$$
$$MIN_MOD(\{\}, P) = \{\{a\}\} \text{ and } EWFM(P) = \{a\}$$

Note this view identifies the intended meaning of rule $a \leftarrow\sim a$ as the equivalent logic formula $a \leftarrow \neg a$, i.e. a. The O-Model of P is empty.

The main differences between ours and their approach are that:

- Like WFS and unlike EWFM, we insist on the supportedness of positive literals, i.e.:

$$\text{An atom } A \in M_P \text{ iff } \exists\ A \leftarrow Body \ \mid\ Body \subseteq M_P$$

- Unlike WFS and unlike EWFM, we relax, by allowing undefined bodies with false heads under certain conditions, the requirement of supportedness of negative literals, i.e. we relax:

$$\sim A \in M_P \text{ iff } \forall\ A \leftarrow Body \ \mid\ Body \text{ is false in } M_P$$

Example 12 Let P be:

$$\begin{array}{lll} c & \leftarrow & \sim b \\ b & \leftarrow & \sim a \\ a & \leftarrow & \sim a \end{array}$$

The O–Model of P is $\{c, \sim b\}$. Note that c has a rule whose body is true in the O–Model and it is not the case that all rules for b are false in it.

The *EWFM* is $\{a, c, \sim b\}$. The atom a is true in the *EWFM* and has no rule with a body true in it. All rules for b have a false body in the *EWFM*. □

Another example where O–Semantics differs from EWFM is the game example of the introduction. In this example EWFM gives the (strange) result that a is a winning position, and thus bets are raised.

A similar approach based on the notion of stable negative hypotheses (built upon the notion of consistency) is introduced in [4], identifying a stable theory associated with a program P as a *"sceptical"* semantics for P, that always contains the well founded model.

One example showing that their approach is still conservative is:

$$\begin{array}{lcl} p & \leftarrow & \sim q \\ q & \leftarrow & \sim r \\ r & \leftarrow & \sim p \\ s & \leftarrow & p \end{array}$$

Stable theories identify the empty set as the meaning of the program; however its O-Model is $\{\sim s\}$, since it is consistent, maximal, sustainable and tenable. Kakas and Mancarella (personal communication) now also obtain this model, as a result of the investigation mentioned in the conclusions of [4]. In this recent work, instead of our notion of sustainable they present:

A is **coherent** if all sets of assumptions B that defeat A are defeated by A.

i.e., if one insists on the set of assumptions A, no consistent evidence to the contrary can be found. No preferred unique model is identified in their approach.

The differences between O–Semantics and Stable Theories are along the same lines as the difference between WFS and Stable Models.

Example 13 Consider program P :

$$\begin{array}{lcl} p & \leftarrow & a \\ p & \leftarrow & b \\ q & \leftarrow & a, b \end{array} \qquad \begin{array}{lcl} a & \leftarrow & \sim b \\ b & \leftarrow & \sim a \end{array}$$

Here $\{\sim a\}$ defeats $\{\sim b\}$ and vice-versa; $\{\sim q\}$ is not defeated by any other; $\{\sim p\}$ is defeated by $\{\sim b\}$ and by $\{\sim a\}$. Thus the O–Model is $\{\sim q\}$.

Both $\{\sim a\}$ and $\{\sim b\}$ are coherent. Thus they identify two models:

$$\{p, a, \sim b, \sim q\} \qquad \{p, b, \sim a, \sim q\}$$

Their intersection is $\{p, \sim q\}$.

As in Stable Models [2], they envisage $\{a \leftarrow\!\sim b;\ b \leftarrow\!\sim a\}$ as $a \vee b$. This is not coherent because then we would expect the special case $a \leftarrow\!\sim a$ to be construed as a, which they don't. □

8 Conclusions

We identify the meaning of a program P as a suitable partial closure of the well founded model of the program in the sense that it contains the well founded model (and thus always exists). The extension we propose reduces undefinedness (which some authors argue is a desirable property) in the intended meaning of a program P, by an adequate form of CWA based on notions of consistency, sustainability and tenability with regard to alternative negative assumptions. Sustainability of a consistent set of negative assumptions insists that there be no other consistent set that defeats it (i.e. there is no hypothetical evidence whose consequences contradict the sustained assumptions). Tenability requires that a maximal sustainable set of assumptions be not contradicted by the consequences of adding to it another competing (nondefeating and nondefeated) maximal sustainable set.

Acknowledgements

We thank ESPRIT BRA COMPULOG (no. 3012), Instituto Nacional de Investigação Científica, Junta Nacional de Investigação Científica e Tecnológica and Gabinete de Filosofia do Conhecimento for their support. We are indebted to the pioneering work of Phan Minh Dung on admissibility of negative hypotheses in Logic Programming, and to Anthony Kakas and Paolo Mancarella for their previous incursions and intuitions into a similar problem in the setting of their Stable Theories. Luís Monteiro is thanked for helpful discussions. Thanks also to an anonymous referee for valuable clarifications.

References

[1] J. Dix. Classifying semantics of logic programs. In A. Nerode, W. Marek, and V. S. Subrahmanian, editors, *Logic Programming and NonMonotonic Reasoning*, pages 166–180. MIT Press, 1991.

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

[3] A. C. Kakas and P. Mancarella. Negation as stable hypothesis. In A. Nerode, W. Marek, and V. S. Subrahmanian, editors, *Logic Programming and NonMonotonic Reasoning*, pages 275–288. MIT Press, 1991.

[4] A. C. Kakas and P. Mancarella. Stable theories for logic programs. In Ueda and Saraswat, editors, *International Logic Programming Symposium*, pages 85–100. MIT Press, 1991.

[5] J. Minker. On indefinite databases and the closed world assumption. In M. Ginsberg, editor, *Readings in Nonmonotonic Reasoning*, pages 326–333. Morgan Kaufmann, 1987.

[6] L. Monteiro. Notes on the negation in logic programs. Technical report, Department of Computer Science-Univerdade Nova de Lisboa, 1992. Course Notes, 3rd Advanced School on Artificial Intelligence, Azores, Portugal, 1992.

[7] L. M. Pereira, J. J. Alferes, and J. N. Aparício. Contradiction Removal within Well Founded Semantics. In A. Nerode, W. Marek, and V. S. Subrahmanian, editors, *Logic Programming and NonMonotonic Reasoning*, pages 105–119. MIT Press, 1991.

[8] L. M. Pereira, J. J. Alferes, and J. N. Aparício. The extended stable models of contradiction removal semantics. In P. Barahona, L. M. Pereira, and A. Porto, editors, *5th Portuguese AI Conference*, volume 541 of *LNAI*, pages 105–119. Springer-Verlag, 1991.

[9] L. M. Pereira, J. N. Aparício, and J. J. Alferes. Counterfactual reasoning based on revising assumptions. In Ueda and Saraswat, editors, *International Logic Programming Symposium*, pages 566–577. MIT Press, 1991.

[10] L. M. Pereira, J. N. Aparício, and J. J. Alferes. Hypothetical reasoning with well founded semantics. In B. Mayoh, editor, *Scandinavian Conference on AI*. IOS Press, 1991.

[11] L. M. Pereira, J. N. Aparício, and J. J. Alferes. Nonmonotonic reasoning with well founded semantics. In Koichi Furukawa, editor, *8th Int. Conf. on Logic Programming*, pages 475–489. MIT Press, 1991.

[12] H. Przymusinska and T. Przymusinski. Semantic issues in deductive databases and logic programs. In R. Banerji, editor, *Formal Techniques in Artificial Intelligence*, pages 321–367. North Holland, 1990.

[13] T. Przymusinski. Extended stable semantics for normal and disjunctive programs. In Warren and Szeredi, editors, *7th Int. Conf. on Logic Programming*, pages 459–477. MIT Press, 1990.

[14] A. Van Gelder, K. A. Ross, and J. S. Schlipf. The well-founded semantics for general logic programs. *Journal of ACM*, 38(3):620–650, 1991.

A Support Sets

In this section we recall the definition and some properties of support sets, introduced in [7].

Definition A.1 (Support Set) *A Support Set of a literal L belonging to the WF Model M_P of a program P, represented as $SS_P(L)$, or $SS(L)$ for short, is obtained as follows:*

- *If L is an atom:*
  - *Choose some rule of P for L where all the literals in its body belong to M_P. One $SS(L)$ is obtained by taking all those body literals plus the literals in some SS of each body literal.*
- *If $L = \sim A$:*
  - *If there are no rules for A in P then the only $SS(L)$ is $\{\}$.*
  - *Otherwise, choose from each rule defined for A, a literal such that its complement belongs to M_P. A $SS(L)$ has all those complement literals, and the literals of a SS of each of them.*

By considering all possible rules of P for a literal all its SSs are obtained.

Here we define $Rules(SS_P(L)) \subseteq P$ as the rules used in the definition above to build $SS_P(L)$.

Proposition A.1 (Existence of Support Set) *Every literal L belonging to the WF Model of a program P has at least one support set $SS_P(L)$.*

Since by definition every literal L with a support set $SS_P(L)$ belong to the WFM of P, we can say that a literal has at least one support set iff it belongs to the WFM.

Other properties of support sets, which are used in some proofs of this paper, are presented below.

Proposition A.2 *For any atom A such that $A \in WFM(P)$, there is at least one support set S of A such that $A \notin S$ and $\sim A \notin S$.*

Proposition A.3 *Let P be a program, $A \in WFM(P)$ be an atom, and $SS_P(A)$ a support set of A. Then $A \in WFM(P')$ for every program P' such that:*

$$Rules(SS_P(A)) \subseteq P' \subseteq P$$