Foundations of Data and Knowledge Systems

VU 181.212, WS 2009

5. Declarative Semantics of Rule Languages

Thomas Eiter and Reinhard Pichler

Institut für Informationssysteme
Technische Universität Wien

November 30, 2010

Outline

Outline

Minimal Model Semantics of Definite Rules

Recall

- Definite programs are finite sets of definite clauses, also called definite rules: $A \leftarrow B_1 \wedge \ldots \wedge B_n$ with $n \geq 0$.
- Definite programs admit a very natural semantics definition:
  - Each program Π is satisfiable.
  - The intersection of all its Herbrand models is a model of Π.
  - This is the *minimal model* of Π.
  - Precisely the atoms implied by Π are true in the minimal model.
- Definite rules are a special case of universal and inductive formulas.
- The interesting model-theoretic properties of definite rules are inherited from these more general classes of formulas.

Theorem

Each set S of definite rules (i.e., each definite program) has a unique minimal Herbrand model. This model is the intersection of all Herbrand models of S. It satisfies precisely those ground atoms that are logical consequences of S.

Minimal Models beyond Herbrand Interpretations

Generalisation

- Minimal Models are also defined for non-Herbrand interpretations
- They make sense also for generalizations of non-inductive formulas
- Uniqueness and intersection property might be lost
- Still the results can be useful

Definition (Generalised Rules)

A generalised rule is a formula of the form $\forall^*(\psi \leftarrow \varphi)$ where φ is positive and ψ is positive and quantifier-free.

Example

The rule $(p(a) \vee p(b) \leftarrow \top)$ is a generalised rule (which is indefinite).

Generalised rules are not necessarily universal: $p(a) \leftarrow \forall x.q(x)$

Supportedness in Minimal Models

Definition (Supported Atoms)

Let $\mathcal{I}$ be an interpretation, V a variable assignment in $dom(\mathcal{I})$ and $A = p(t_1, \ldots, t_n)$ an atom, $n \geq 0$.

- an atom B supports A in $\mathcal{I}[V]$ iff $\mathcal{I}[V] \models B$ and $B = p(s_1, \ldots, s_n)$ and $s_i^{\mathcal{I}[V]} = t_i^{\mathcal{I}[V]}$ for $1 \leq i \leq n$.
- a set C of atoms supports A in $\mathcal{I}[V]$ iff $\mathcal{I}[V] \models C$ and there is an atom in C that supports A in $\mathcal{I}[V]$.
- a generalised rule $\forall^*(\psi \leftarrow \varphi)$ supports A in I iff for each variable assignment V with $\mathcal{I}[V] \models \varphi$ there is an implicant C of ψ that supports A in $\mathcal{I}[V]$.

Informally, an implicant C of ψ is a set of atoms which logically implies ψ

Implicant of a Positive Quantifier-Free Formula

Definition (Pre-Implicant and Implicant)

Let ψ be a positive quantifier-free formula. The set $primps(\psi)$ of pre-implicants of ψ is defined as follows:

- $primps(\psi) = \{ \{\psi\} \}$ if ψ is an atom or $\top$ or $\bot$.
- $primps(\neg\psi_1) = primps(\psi_1)$.
- $primps(\psi_1 \wedge \psi_2) = \{ C_1 \cup C_2 \mid C_1 \in primps(\psi_1),\ C_2 \in primps(\psi_2) \}$.
- $primps(\psi_1 \vee \psi_2) = primps(\psi_1 \Rightarrow \psi_2) = primps(\psi_1) \cup primps(\psi_2)$.

The set of implicants of ψ is obtained from $primps(\psi)$ by removing all sets containing $\bot$ and by removing $\top$ from the remaining sets.

Lemma

1. *If C is an implicant of ψ, then $C \models \psi$.*
2. *For any interpretation $\mathcal{I}$, if $\mathcal{I} \models \psi$ then there exists an implicant C of ψ with $\mathcal{I} \models C$.*

Supportedness Result

Theorem (Minimal Models Satisfy Only Supported Ground Atom)

Let S be a set of generalised rules. If $\mathcal{I}$ is a minimal model of S, then for each ground atom A with $\mathcal{I} \models A$ there is a generalised rule in S that supports A in $\mathcal{I}$.

Example

Consider a signature containing a unary relation symbol p and constants a and b.
Let $S = \{ (p(b) \leftarrow \top) \}$.
The interpretation $\mathcal{I}$ with $dom(\mathcal{I}) = \{1\}$ and $a^{\mathcal{I}} = b^{\mathcal{I}} = 1$ and $p^{\mathcal{I}} = \{(1)\}$ is a minimal model of S.
Moreover, $\mathcal{I} \models p(a)$. By the theorem, $p(a)$ is supported in $\mathcal{I}$ by $p(b)$, which can be confirmed by applying the definition.

Non-Minimal Supportedness

The converse of the Theorem fails, e.g. $S = \{ (p \leftarrow p) \}$.

Proof

Assume that $\mathcal{I}$ is a minimal model of S with domain D and there is a ground atom A with $\mathcal{I} \models A$, such that no $r \in S$ supports A in $\mathcal{I}$.

Let $\mathcal{I}'$ be identical to $\mathcal{I}$ except that $\mathcal{I}' \not\models A$. Then $\mathcal{I}' < \mathcal{I}$.

Consider any $r = \forall^*(\psi \leftarrow \varphi)$ from S. By assumption, r does not support A. Let V be an arbitrary variable assignment in D. We show $\mathcal{I}'[V] \models (\psi \leftarrow \varphi)$.

If $\mathcal{I}[V] \not\models \varphi$, as ψ is positive, also $\mathcal{I}'[V] \not\models \varphi$; hence $\mathcal{I}'[V] \models (\psi \leftarrow \varphi)$.

If $\mathcal{I}[V] \models \varphi$, then $\mathcal{I}[V] \models \psi$ because $\mathcal{I}$ is a model of S. Furthermore, by assumption for each implicant C of ψ either $\mathcal{I}[V] \not\models C$ or no atom in C supports A in $\mathcal{I}[V]$. Consider two cases:

- If $\mathcal{I}[V] \not\models C$ for each implicant C of ψ, then $\mathcal{I}[V] \not\models \psi$ by the above Lemma (part 2); contradiction.
- If $\mathcal{I}[V] \models C$ for some implicant C of ψ, then by assumption no atom in C supports A in $\mathcal{I}[V]$. By construction, $\mathcal{I}'[V]$ agrees with $\mathcal{I}[V]$ on all atoms except those supporting A in $\mathcal{I}[V]$, thus $\mathcal{I}'[V] \models C$. By the above Lemma (part 1), $\mathcal{I}'[V] \models \psi$. Hence $\mathcal{I}'[V] \models (\psi \leftarrow \varphi)$.

In all possible cases $\mathcal{I}'$ satisfies r; thus $\mathcal{I}'$ is a model of S, contradicting the minimality of $\mathcal{I}$. □

Semantic vs Syntactic Support

- The above theorem is semantic in nature:
  In the above example, $p(a)$ is supported by $p(b)$
- There is no syntactic connection between these atoms.
- It holds under suitable conditions.

Definition (Unique Name Property)

An interpretation $\mathcal{I}$ has the *unique name property*, if for each term s, ground term t, and variable assignment V in $dom(\mathcal{I})$ with $s^{\mathcal{I}[V]} = t^{\mathcal{I}[V]}$ there exists a substitution σ with $s\sigma = t$.

- Herbrand interpretations have the unique name property.
- The relationship between the supporting atom and the supported ground atom specialises to the (syntactic and decidable) ground instance relationship.
- Sometimes, unique names are postulated (*Unique Names Assumption*)

Outline

Minimal Model Construction

Outline

- The minimal models semantics is not constructive.
- We need algorithms to compute the / reason from the minimal model
- Different methods exist, including
  - algebraic approaches (fixpoints of consequence operators, "bottom up")
  - proof-theoretic approaches (special resolution procedures, "top down")
- We consider here first fix-point construction, for which we need concepts from operator theory.
- We confine here to a specific case of operators, applied to elements M of the powerset $\mathcal{P}(X)$ (the set of subsets) of a set X.

Operators

Definition (Operator)

Let X be a set. An operator on X is a mapping $\Gamma : \mathcal{P}(X) \to \mathcal{P}(X)$.

Definition (Monotonic operator)

Let X be a set. An operator Γ on X is *monotonic*, iff for all subset $M \subseteq M' \subseteq X$ holds: $\Gamma(M) \subseteq \Gamma(M')$.

Definition (Continuous operator)

Let X be a nonempty set.
A set $Y \subseteq \mathcal{P}(X)$ of subsets of X is *directed*, if every finite subset of Y has an *upper bound* in Y, i.e., for each finite $Y_{fin} \subseteq Y$, there is a set $M \in Y$ such that $\bigcup Y_{fin} \subseteq M$.
An operator Γ on X is *continuous*, iff for each directed set $Y \subseteq \mathcal{P}(X)$ of subsets of X holds: $\Gamma(\bigcup Y) = \bigcup\{\Gamma(M) \mid M \in Y\}$.

Continuous vs Monotone Operators

Lemma

Each continuous operator on a nonempty set is monotonic.

Proof.

Let Γ be a continuous operator on $X \neq \emptyset$. Let $M \subseteq M' \subseteq X$. Since Γ is continuous, $\Gamma(M') = \Gamma(M \cup M') = \Gamma(M) \cup \Gamma(M')$, thus $\Gamma(M) \subseteq \Gamma(M')$. □

The converse does not hold.

Example

Let $\Gamma(X) = \emptyset$, if X is finite, and $\Gamma(X) = X$, if X is infinite.

- Γ is monotonic.
- Γ is not continuous in general. E.g., let $X = \mathbb{N}$ and $Y = \{\{0, 1, \ldots n\} \mid n \in \mathbb{N}\}$.
  Then $\Gamma(\bigcup Y) = \mathbb{N}$ but $\bigcup_{M \in Y} \Gamma(M) = \emptyset$.

Fixpoints of Monotonic and Continuous Operators

Definition (Fixpoint)

Let Γ be an operator on a set X. A subset $M \subseteq X$ is

- a *pre-fixpoint* of Γ iff $\Gamma(M) \subseteq M$;
- a *fixpoint* of Γ iff $\Gamma(M) = M$.

Theorem (Knaster-Tarski, existence of least and greatest fixpoint)

Let Γ be a monotonic operator on a nonempty set X. Then Γ has a least fixpoint $lfp(\Gamma)$ and a greatest fixpoint $gfp(\Gamma)$ with

$$\begin{aligned} lfp(\Gamma) &= \bigcap\{M \subseteq X \mid \Gamma(M) = M\} = \bigcap\{M \subseteq X \mid \Gamma(M) \subseteq M\}. \\ gfp(\Gamma) &= \bigcup\{M \subseteq X \mid \Gamma(M) = M\} = \bigcup\{M \subseteq X \mid \Gamma(M) \subseteq M\}. \end{aligned}$$

- This is a fundamental result with many applications in Computer Science.
- It holds for more general structures (complete partial orders).

Proof.

For the least fixpoint let $L = \bigcap\{M \subseteq X \mid \Gamma(M) \subseteq M\}$.

Consider an arbitrary $M \subseteq X$ with $\Gamma(M) \subseteq M$.

By definition of L we have $L \subseteq M$. Since Γ is monotonic, $\Gamma(L) \subseteq \Gamma(M)$. With the assumption $\Gamma(M) \subseteq M$ follows $\Gamma(L) \subseteq M$. Therefore

$$\Gamma(L) \subseteq \bigcap\{M \subseteq X \mid \Gamma(M) \subseteq M\} = L. \quad (1)$$

For the opposite inclusion, from (1) and since Γ is monotonic it follows that $\Gamma(\Gamma(L)) \subseteq \Gamma(L)$. By definition of L therefore

$$L \subseteq \Gamma(L). \quad (2)$$

From (1) and (2) it follows that L is a fixpoint of Γ.

Now let $L' = \bigcap\{M \subseteq X \mid \Gamma(M) = M\}$.

Then $L' \subseteq L$, because L is a fixpoint of Γ.

The opposite inclusion $L \subseteq L'$ holds, since every set M involved in the intersection defining L' is also involved in the intersection defining L.

The proof for the greatest fixpoint is similar. □

Ordinal Powers

Ordinal numbers

- Ordinal numbers are the *order types* of *well-ordered sets* (i.e., totally order sets where each set has a minimum.)
- The generalize natural numbers, and can be defined as *hereditarily transitive sets* (J. von Neumann).
- There are *successor ordinals* β, given by $\beta = \alpha + 1$ for ordinal α, and limit ordinals λ (not of this form).
- The first limit ordinal, ω, corresponds to the set $\mathbb{N}$ of all natural numbers.

Definition (Ordinal powers of a monotonic operator)

Let Γ be a monotonic operator on a nonempty set X. For each ordinal, the *upward and downward power* of Γ is defined as

$\Gamma\uparrow 0 = \emptyset$	(base case)	$\Gamma\downarrow 0 = X$
$\Gamma\uparrow \alpha{+}1 = \Gamma(\Gamma\uparrow\alpha)$	(successor case)	$\Gamma\downarrow\alpha{+}1 = \Gamma(\Gamma\downarrow\alpha)$
$\Gamma\uparrow\lambda = \bigcup\{\Gamma\uparrow\beta \mid \beta<\lambda\}$	(limit case)	$\Gamma\downarrow\lambda = \bigcap\{\Gamma\downarrow\beta \mid \beta<\lambda\}$

Lemma

Let Γ be a monotonic operator on a nonempty set X. For each ordinal α holds:

1. $\Gamma\uparrow\alpha \;\subseteq\; \Gamma\uparrow\alpha+1$
2. $\Gamma\uparrow\alpha \;\subseteq\; \mathit{lfp}(\Gamma)$.
3. *If* $\Gamma\uparrow\alpha = \Gamma\uparrow\alpha+1$, *then* $\mathit{lfp}(\Gamma) = \Gamma\uparrow\alpha$.

Idea.

1. and 2. are shown by transfinite induction on α. Item 3. is shown as follows: If $\Gamma\uparrow\alpha = \Gamma\uparrow\alpha+1$, then $\Gamma\uparrow\alpha = \Gamma(\Gamma\uparrow\alpha)$, i.e., $\Gamma\uparrow\alpha$ is a fixpoint of Γ, therefore $\Gamma\uparrow\alpha \subseteq \mathit{lfp}(\Gamma)$ by 2., and $\mathit{lfp}(\Gamma)\subseteq\Gamma\uparrow\alpha$ by definition. □

Theorem

For any monotonic operator Γ on $X\neq\emptyset$, $\mathit{lfp}(\Gamma) = \Gamma\uparrow\alpha$ for some ordinal α.

Proof.

Otherwise, for all ordinals α by the previous lemma $\Gamma\uparrow\alpha\subseteq\Gamma\uparrow\alpha+1$ and $\Gamma\uparrow\alpha\neq\Gamma\uparrow\alpha+1$. Thus $\Gamma\uparrow$ injectively maps the ordinals to $\mathcal{P}(X)$, a contradiction (there are "more" ordinals than any set can have elements). □

Least Fixpoint of Continuous Operator

Theorem (Kleene)

Let Γ be a continuous operator on a nonempty set X. Then

$$\mathit{lfp}(\Gamma) = \Gamma\uparrow\omega.$$

Proof.

By 1. from the previous lemma, it suffices to show that $\Gamma\uparrow\omega+1 = \Gamma\uparrow\omega$.

$\Gamma\uparrow\omega+1 = \Gamma(\Gamma\uparrow\omega)$	by definition, successor case
$= \Gamma\big(\bigcup\{\Gamma\uparrow n \mid n\in\mathbb{N}\}\big)$	by definition, limit case
$= \bigcup\{\Gamma(\Gamma\uparrow n) \mid n\in\mathbb{N}\}$	because Γ is continuous
$= \bigcup\{\Gamma\uparrow n+1 \mid n\in\mathbb{N}\}$	by definition, successor case
$= \Gamma\uparrow\omega$	by definition, base case □

Note: An analogous result for the greatest fixpoint does not hold.

Outline

Immediate Consequence Operator

- We now apply the above results for universal generalized definite rules.
- Here $X = HB$ and a subset M is a set $B \subseteq HB$ of ground atoms.

Definition (Immediate consequence operator)

Let S be a set of universal generalised definite rules. Let $B \subseteq HB$ be a set of ground atoms. The *immediate consequence operator* $\mathbf{T}_S$ *for* S is:

$$\begin{array}{ll} \mathbf{T}_S : \mathcal{P}(HB) \rightarrow \mathcal{P}(HB) & \\ \quad B \mapsto \{A \in HB \mid & \text{there is a ground instance } ((A_1 \wedge \ldots \wedge A_n) \leftarrow \varphi) \\ & \text{of a member of } S \text{ with } HI(B) \models \varphi \text{ and } A = A_i \\ & \text{for some } i \text{ with } 1 \leq i \leq n \quad \} \end{array}$$

Lemma ($\mathbf{T}_S$ is continuous)

Let S be a set of universal generalised definite rules. The immediate consequence operator $\mathbf{T}_S$ is continuous (hence, also monotonic).

Theorem

Let S be a set of universal generalised definite rules. Let $B \subseteq HB$ be a set of ground atoms. Then $HI(B) \models S$ iff $\mathbf{T}_S(B) \subseteq B$.

Proof.

"only if:" Assume $HI(B) \models S$. Let $A \in \mathbf{T}_S(B)$, i.e., $A = A_i$ for some ground instance $((A_1 \wedge \ldots \wedge A_n) \leftarrow \varphi)$ of a member of S with $HI(B) \models \varphi$.

By assumption $HI(B) \models (A_1 \wedge \ldots \wedge A_n)$, hence $HI(B) \models A$, hence $A \in B$ because A is a ground atom.

"if:" Assume $\mathbf{T}_S(B) \subseteq B$. Let $r = ((A_1 \wedge \ldots \wedge A_n) \leftarrow \varphi)$ be a ground instance of a member of S. It suffices to show that $HI(B)$ satisfies r.

- If $HI(B) \not\models \varphi$, it does.
- If $HI(B) \models \varphi$, then $A_1 \in \mathbf{T}_S(B), \ldots, A_n \in \mathbf{T}_S(B)$ by definition of $\mathbf{T}_S$.

  By assumption $A_1 \in B, \ldots, A_n \in B$.

  As all A_i are ground atoms, $HI(B) \models A_1, \ldots, HI(B) \models A_n$. Thus $HI(B)$ satisfies r.

□

Corollary (Fixpoint Characterization of the Least Herbrand Model)

a Let S be a set of universal generalised definite rules. Then

(i) $lfp(\mathbf{T}_S) = \mathbf{T}_S \uparrow \omega = Mod_\cap(S) = \{A \in HB \mid S \models A\}$ *and*

(ii) $HI(lfp(\mathbf{T}_S))$ *is the unique minimal Herbrand model of* S.

Proof.

(i): By the Lemma above, $\mathbf{T}_S$ is a continuous operator on HB, and by Kleene's Theorem, $lfp(\mathbf{T}_S) = \mathbf{T}_S \uparrow \omega$. Note that $Mod_{HB}(S) \neq \emptyset$ (as $HI(HB) \models S$) Now,

$$\begin{array}{lll} lfp(\mathbf{T}_S) & = \bigcap\{B \subseteq HB \mid \mathbf{T}_S(B) \subseteq B\} & \text{by the Knaster-Tarski Theorem} \\ & = \bigcap\{B \subseteq HB \mid HI(B) \models S\} & \text{by the previous Theorem} \\ & = \bigcap Mod_{HB}(S) & \text{by definition of } Mod_{HB} \\ & = Mod_\cap(S) & \text{by definition of } Mod_\cap \\ & = \{A \in HB \mid S \models A\} & \text{as } S \text{ is universal (see unit 4)} \end{array}$$

(ii): By (i), $HI(lfp(\mathbf{T}_S))$ is the intersection of all Herbrand models of S, and $HI(lfp(\mathbf{T}_S)) \models S$, as S is satisfiable.

Hence, $HI(lfp(\mathbf{T}_S))$ is the unique minimal Herbrand model of S. □

Charcterization Summary

- The "natural meaning" of a set S of universal generalised definite rules can defined in different but equivalent ways:
  - as the unique minimal Herbrand model of S;
  - as the intersection $HI(Mod_\cap(S))$ of all Herbrand models of S;
  - as the set $\{A \in HB \mid S \models A\}$ of ground atoms entailed by S;
  - as the least fixpoint $lfp(\mathbf{T}_S)$ of the immediate consequence operator
- Declarative and procedural (forward chaining) semantics coincide.
- Further equivalent procedural semantics, based on SLD resolution, exists (backward chaining).

Outline

Declarative Semantics of Rules with Negation

If a database of students does not list "Mary", then one may conclude that "Mary" is not a student. The principle underlying this is called closed world assumption (CWA).

Two approaches to coping with this form of negation:

- axiomatization within first-oder predicate logic
- deduction methods not requiring specific axioms conveying the CWA

The second approach is desirable but it poses the problem of the declarative semantics, or model theory.

Not all Minimal Models convey the CWA

Example

- $S_1 = \{ (q \leftarrow r \wedge \neg p),\ (r \leftarrow s \wedge \neg t),\ (s \leftarrow \top) \}$
  Minimal Herbrand models: $HI(\{s, r, q\})$, $HI(\{s, r, p\})$, and $HI(\{s, t\})$.
  Intuitively, p and r are not "justified" by the rules on S_1.
- $S_2 = \{ (p \leftarrow \neg q),\ (q \leftarrow \neg p) \}$
  Minimal Herbrand models: $HI(\{p\})$, $HI(\{q\})$.
  Intuitively, exactly one of p and q should be true, but it is unclear which.
- $S_3 = \{ (p \leftarrow \neg p) \}$
  Minimal Herbrand model: $HI(\{p\})$.
  p can not be arguably justified from S_3, which is intuitively not consistent.
- $S_4 = \{ (p \leftarrow \neg p),\ (p \leftarrow \top) \}$
  Minimal Herbrand model: $HI(\{p\})$.
  Here, p is arguably justified and S_4 should be consistent.

  Note: different from classical logic, a subset of a consistent rule set ($S_3 \subseteq S_4$) may be inconsistent!

Non-Monotonic Consequence

- A consequence operator is a mapping that assigns a set S of formulas a set of formulas $Th(S)$ (satisfying certain properties).
- We can view $Th(S)$ as an operator considered above.
- S_3 and S_4 suggest that a consequence operator for rules with negation should be non-monotonic (if $Th(S)$ for "inconsistent" S yields all formulas).
- But also for "consistent" sets of formulas, consequence should act non-monotonic, if it is based on *canonical models*, which are *preferred minimal Herbrand models* (denoted $Th_{can}(S)$).

Example

$S_5 = \{ (q \leftarrow \neg p) \}$ has the minimal Herbrand models: $HI(\{p\})$ and $HI(\{q\})$.
Only $HI(\{q\})$ conveys the intuitive meaning under the CWA and should be retained as (the only) canonical model. Therefore, $q \in Th_{can}(S_5)$.

$S_5' = S_5 \cup \{ (p \leftarrow \top) \}$ has the single minimal Herbrand model $HI(\{p\})$, which also conveys the intuitive meaning under the CWA and should be retained as a canonical model. Therefore, $q \notin Th_{can}(S_5')$.

Thus, $S_5 \subseteq S_5'$, but $Th_{can}(S_5) \not\subseteq Th_{can}(S_5')$.

Outline

Stratifiable Rule Sets

Basic Idea

Avoid cases like $(p \leftarrow \neg p)$ and more generally recursion through negative literals.

Definition (Stratification)

Let S be a set of normal clauses (rules). A stratification of S is a partition $S_0, \ldots, S_k$ of S such that

- For each relation symbol p there is a stratum S_i, such that all clauses of S containing p in their consequent are members of S_i.
  In this case one says that the relation symbol p *is defined in stratum* S_i.
- For each stratum S_j and positive literal A in the antecedents of members of S_j, the relation symbol of A is defined in a stratum S_i with $i \leq j$.
- For each stratum S_j and negative literal $\neg A$ in the antecedents of members of S_j, the relation symbol of A is defined in a stratum S_i with $i < j$.

A set of normal clauses is called stratifiable, if there exists a stratification of it.

Example

- Each definite program is stratifiable by making it its only stratum.
- The set $S = \{ (r \leftarrow \top), (q \leftarrow r), (p \leftarrow q \wedge \neg r) \}$ is stratifiable:
  the stratum S_0 contains the first clause and the stratum S_1 the last one, while the middle clause may belong to either of the strata.
- The set $S = \{ (p \leftarrow \neg p) \}$ is not stratifiable.
- Any set of normal clauses with a "cycle of recursion through negation" (defined syntactically via a *dependency graph*) is not stratifiable.

Stratifiable Rule Sets – Canoncial Model

Principal Idea

- The stratum S_0 always consists of definite clauses (positive definite rules).
- Hence the truth values of all atoms of stratum S_0 can be determined without negation being involved.
- After that the clauses of stratum S_1 refer only to such negative literals whose truth values have already been determined in S_0.
- After that the clauses of stratum S_2 refer only to such negative literals whose truth values have already been determined in S_0 and S_1.
- And so on.

That is, *work stratum by stratum*.

Stratification Theorem (Apt, Blair and Walker)

Each stratifiable rule set has a well-defined canonical model (also called *perfect model*), which is *independent of a particular stratification*.

Outline

Stable Model Semantics

Basic Idea

Perform *assumption-based* evaluation, where negation takes the value in the final result.

Definition (Gelfond-Lifschitz transformation)

Let S be a (possibly infinite) set of *ground* normal clauses, i.e., of formulas

$$A \leftarrow L_1 \wedge \ldots \wedge L_n$$

where $n \geq 0$ and A is a ground atom and the L_i for $1 \leq i \leq n$ are ground literals. Let $B \subseteq HB$. The Gelfond-Lifschitz transform $GL_B(S)$ of S with respect to B is obtained from S as follows:

1. remove each clause whose antecedent contains a literal $\neg A$ with $A \in B$.
2. remove from the antecedents of the remaining clauses all negative literals.

Definition (Stable model)

Let S be a (possibly infinite) set of ground normal clauses. An Herbrand interpretation $HI(B)$ is a *stable model of S* iff it is the unique minimal Herbrand model of $GL_B(S)$.
A stable model of a set S of normal clauses is a stable model of the (possibly infinite) set of ground instances of S.

Example

- $S_1 = \{\, (q \leftarrow r \wedge \neg p),\ (r \leftarrow s \wedge \neg t),\ (s \leftarrow \top)\,\}$ has one stable model: $HI(\{s, r, q\})$.
- $S_2 = \{\, (p \leftarrow \neg q),\ (q \leftarrow \neg p)\,\}$ has two stable models: $HI(\{p\})$ and $HI(\{q\})$.
- $S_3 = \{\, (p \leftarrow \neg p)\,\}$ has no stable model.
- $S_4 = \{\, (p \leftarrow \neg p),\ (p \leftarrow \top)\,\}$ has one stable model: $HI(\{p\})$.

Stable Model Semantics –Properties

Theorem

Each stable model of a normal clause set S is a minimal Herbrand model of S.

Proof.

It suffices to consider a set S of *ground* normal clauses.
As easily seen, $HI(B) \models GL_B(S)$ implies $HI(B) \models S$.
Let $B' \subseteq B \subseteq HB$ such that $HI(B)$ is a stable model of S and $HI(B')$ is also a model of S, i.e., $HI(B') \models S$. If we establish that $HI(B') \models GL_B(S)$, then $B' = B$ by the minimality of a stable model.
Let $C \in GL_B(S)$. By definition of $GL_B(S)$ there exists a clause $D \in S$, such that C is obtained from D by removing the negative literals from its antecedent. If $\neg A$ is such a literal, then $A \notin B$, and, since $B' \subseteq B$, also $A \notin B'$. Therefore, $C \in GL_{B'}(S)$. As $HI(B') \models S$, it follows $HI(B') \models C$. □

Proposition

Each stratifiable rule set has exactly one stable model, which coincides with the respective canonical model.

Stable Model Semantics – Evaluation

- The Stable Model Semantics coincides with the intuitive understanding based on the "Justification Postulate".
- It does not satisfy the "Consistency Postulate".
- It gracefully generalizes the canonical semantics.
- To date, Stable Model Semantics is the predominant multiple model non-montonic semantics for rule sets with negation.

Outline

Well-Founded Semantics

Basic Idea

- Avoid cases like $(p \leftarrow \neg p)$ by using a third truth value, *unkown*.
- Try to build a single *partial* model, in which p would be unknown.

Notation

For a literal L, $\overline{L}$ is its complement with $\overline{A} = \neg A$ and $\overline{\neg A} = A$ for an atom A.
For a set I of ground literals,

$$\overline{I} = \{\, \overline{L} \mid L \in I \,\}, \quad pos(I) = I \cap HB, \quad neg(I) = \overline{I} \cap HB.$$

Thus, $I = pos(I) \cup \overline{neg(I)}$.

Definition

A set I of ground literals is *consistent*, iff $pos(I) \cap neg(I) = \emptyset$. Otherwise, I is *inconsistent*.
Two sets I_1 and I_2 of ground literals are *(in)consistent* iff $I_1 \cup I_2$ is.
A literal L and a set I of ground literals are *(in)consistent* iff $\{L\} \cup I$ is.

Definition (Partial interpretation)

A *partial interpretation* is a consistent set I of ground literals; it is *total*, iff $pos(I) \cup neg(I) = HB$, i.e., for each ground atom A either $A \in I$ or $\neg A \in I$.
For a total I, the Herbrand interpretation induced by I is $HI(I) = HI(pos(I))$.

Definition (Satisfaction for partial interpretations)

Let I be a partial interpretation.

Then $\top$ is *satisfied* in I and $\bot$ is *falsified* in I.

A ground literal L is
satisfied or true in I iff $L \in I$.
falsified or false in I iff $\overline{L} \in I$.
undefined in I iff $L \notin I$ and $\overline{L} \notin I$.

A conjunction $L_1 \wedge \ldots \wedge L_n$ of ground literals, $n \geq 0$, is
satisfied or true in I iff each L_i for $1 \leq i \leq n$ is satisfied in I.
falsified or false in I iff at least one L_i for $1 \leq i \leq n$ is falsified in I.
undefined in I iff each L_i for $1 \leq i \leq n$ is satisfied or undefined in I and at least one of them is undefined in I.

Definition (Satisfaction, cont'd)

Let I be a partial interpretation.
A ground normal clause $A \leftarrow \varphi$ is
satisfied or true in I iff A is satisfied in I or φ is falsified in I.
falsified or false in I iff A is falsified in I and φ is satisfied in I.
weakly falsified in I iff A is falsified in I and φ is satisfied or undefined in I.

A normal clause is
satisfied or true in I iff each of its ground instances is.
falsified or false in I iff at least one of its ground instances is.
weakly falsified in I iff at least one of its ground instances is.

A set of normal clauses is
satisfied or true in I iff each of its members is.
falsified or false in I iff at least one of its members is.
weakly falsified in I iff at least one of its members is.

- Note: "weakly falsified" intuitively means that by turning from "undefined" to "true", the clause could be falsified.
- For a total interpretation I, the cases "undefined" and "weakly falsified" are impossible, and satisfaction in $HI(I)$ amounts to the classical notion.

Definition (Total and partial model)

Let S be a set of normal clauses.
A total interpretation I is a *total model of* S, iff S is satisfied in I.
A partial interpretation I is a *partial model of* S, iff there exists a total model I' of S with $I \subseteq I'$.

- If a ground normal clause C is weakly falsified, but not falsified in a partial interpretation I, then its consequent *is* falsified in I and some literal L in its antecedent are undefined in I.
- No extension of I with additional literals can satisfy the consequent.
- The only way to satisfy S is to extend I by the complement $\overline{L}$ of some undefined antecedent literal L (which falsifies the antecedent).
- Any extension of I that satisfies all antecedent literals L falsifies C.

Lemma (Weak Falsification)

Let S be a set of normal clauses and I a partial interpretation. If no clause in S is weakly falsified in I, then I is a partial model of S.

Unfounded Sets

Principle for Drawing Negative Conclusions

Given a partial interpretation I, a set U of ground atoms is "unfounded" wrt a clause set, if each atom A in U is unjustified wrt I, *taking U into account*.

Example

Let $S = \{(p \leftarrow q),\ (q \leftarrow p)\}$. For $U = \{p, q\}$, p, q are unjustified wrt $\{p, q\}$.

Definition (Unfounded set of ground atoms)

Let S be a set of normal clauses, and I a partial interpretation.
A set $U \subseteq HB$ of ground atoms is an *unfounded set* wrt S and I, if for each $A \in U$ and for each ground instance $r = A \leftarrow L_1 \wedge \ldots \wedge L_n$, $n \geq 1$, of a member of S, at least one of the following holds:

1. $L_i \in \overline{I}$ for some positive or negative L_i with $1 \leq i \leq n$. (L_i is falsified in I)
2. $L_i \in U$ for some positive L_i with $1 \leq i \leq n$. (L_i is unfounded)

A respective L_i is a *witness of unusability* for r.

Example

- Let $S = \{(p \leftarrow q),\ (q \leftarrow p)\}$.
  Then $U = \{p, q\}$ is an unfounded set wrt S and $I = \{p, q\}$.
  Both a and b are unfounded by condition 2.
- Let $S' = \{\ (q \leftarrow p),\ (r \leftarrow s),\ (s \leftarrow r)\ \}$ and $I = \{\neg p, \neg q\}$.
  The set $U' = \{q, r, s\}$ is unfounded wrt S' and I.
  The atom q is unfounded by condition 1, the atoms r and s by condition 2.

Lemma

Let S be a set of normal clauses and I a partial interpretation. There exists a unique maximal (under set inclusion) unfounded set with respect to S and I, $GUS_S(I)$, which is the union of all unfounded sets with respect to S and I.

Example (cont'd)

$GUS_S(I) = \{p, q\}$ and $GUS_{S'}(I') = \{p, q, r, s\}$

Observation

- If all atoms in I are founded, by switching any unfounded atom(s) all rules remain satisfied.
- As no backtracking is needed, unfounded atoms can be safely made false.

Lemma

Let S be a set of normal clauses, I be a partial interpretation, and U' be an unfounded set with respect to S and I, such that $pos(I) \cap U' = \emptyset$.
For each $U \subseteq U'$, its remainder $U' \setminus U$ is unfounded w.r.t. S and $I \cup \overline{U}$.

A kind of opposite property is that false atoms are unfounded.

Lemma

Let S be a set of normal clauses and $I = pos(I) \cup \overline{neg(I)}$ be a partial interpretation. If no clause in S is weakly falsified in I, then $neg(I)$ is unfounded with respect to S and $pos(I)$.

The above properties are exploited to extend a partial interpretation.

Definition (Operators $\mathbf{T}_S$, $\mathbf{U}_S$, $\mathbf{W}_S$)

Let $\mathcal{PI} = \{\, I \subseteq HB \cup \overline{HB} \mid I \text{ is consistent} \,\}$, and note that $\mathcal{P}(HB) \subseteq \mathcal{PI}$.
Let S be a set of normal clauses. We define three operators:

$$\begin{array}{llll}
\mathbf{T}_S: & \mathcal{PI} & \to & \mathcal{P}(HB) \\
 & I & \mapsto & \{\, A \in HB \mid \text{there is a ground instance } (A \leftarrow \varphi) \\
 & & & \quad \text{of a member of } S \text{ such that } \varphi \text{ is satisfied in } I \,\} \\
\mathbf{U}_S: & \mathcal{PI} & \to & \mathcal{P}(HB) \\
 & I & \mapsto & \text{the maximal subset of } HB \text{ that is unfounded wrt } S \text{ and } I \\
\mathbf{W}_S: & \mathcal{PI} & \to & \mathcal{PI} \\
 & I & \mapsto & \mathbf{T}_S(I) \cup \overline{\mathbf{U}_S(I)}
\end{array}$$

- Starting from "knowing" I, the ground atoms in $\mathbf{T}_S(I)$ have to be true;
- those in $\mathbf{U}_S(I)$ are unfounded;
- $\mathbf{T}_S(I) \cap \mathbf{U}_S(I) = \emptyset$, thus $\mathbf{W}_S(I)$ is consistent.

Lemma

$\mathbf{T}_S$, $\mathbf{U}_S$, and $\mathbf{W}_S$ are monotonic.

Example

Suppose $HB = \{p, q, r, s, t\}$, and let $I_0 = \emptyset$ and
$S = \{\, (q \leftarrow r \wedge \neg p),\ (r \leftarrow s \wedge \neg t),\ (s \leftarrow \top) \,\}$.

$$\begin{array}{lcl}
\mathbf{T}_S(I_0) & = & \{s\} \\
\mathbf{U}_S(I_0) & = & \{p, t\} \\
\mathbf{W}_S(I_0) & = & \{s, \neg p, \neg t\} = I_1 \\[1ex]
\mathbf{T}_S(I_1) & = & \{s, r\} \\
\mathbf{U}_S(I_1) & = & \{p, t\} \\
\mathbf{W}_S(I_1) & = & \{s, r, \neg p, \neg t\} = I_2 \\[1ex]
\mathbf{T}_S(I_2) & = & \{s, r, q\} \\
\mathbf{U}_S(I_2) & = & \{p, t\} \\
\mathbf{W}_S(I_2) & = & \{s, r, q, \neg p, \neg t\} = I_3 \\[1ex]
\mathbf{T}_S(I_3) & = & \{s, r, q\} \\
\mathbf{U}_S(I_3) & = & \{p, t\} \\
\mathbf{W}_S(I_3) & = & I_3
\end{array}$$

Theorem (Existence of least fixpoint)

Let S be a set of normal clauses. (1) The operator $\mathbf{W}_S$ has a least fixpoint given by $\mathit{lfp}(\mathbf{W}_S) = \bigcap\{I \in \mathcal{PI} \mid \mathbf{W}_S(I) = I\} = \bigcap\{I \in \mathcal{PI} \mid \mathbf{W}_S(I) \subseteq I\}$. Moreover, (2) $\mathit{lfp}(\mathbf{W}_S)$ is a partial interpretation of S and (3) $\mathit{lfp}(\mathbf{W}_S)$ is a partial model of S.

Proof.

Part one follows from the Knaster-Tarski Theorem. For part two, both consistency and that no clause in S is weakly falsified, are shown by *transfinite* induction. The Weak Falsification Lemma ensures the model property. □

Definition (Well-founded model)

The well-founded model of a set S of normal clauses is $\mathit{lfp}(\mathbf{W}_S)$.

- The well-founded model may be total (it specifies a truth value for each ground atom) or partial (it leaves some atoms undefined).
- If S is stratifiable, then S has a total well-founded model, which coincides with the canonical (perfect model).

Example

- $S_1 = \{ (q \leftarrow r \wedge \neg p),\ (r \leftarrow s \wedge \neg t),\ (s \leftarrow \top) \}$ has the well-founded model $\{s, r, q, \neg p, \neg t\}$. It is total.
- $S_2 = \{ (p \leftarrow \neg q),\ (q \leftarrow \neg p) \}$ has the well-founded model $\emptyset$. It is partial and leaves the truth values of p and of q undefined.
- $S_3 = \{ (p \leftarrow \neg p) \}$ has the well-founded model $\emptyset$. It is partial and leaves the truth value of p undefined.
- $S_4 = \{ (p \leftarrow \neg p),\ (p \leftarrow \top) \}$ has the well-founded model $\{p\}$. It is total.

Well-Founded Semantics - Evaluation

- The well-founded semantics coincides with an intuitive understanding based on the "Justification Postulate".
- A set of normal clauses always has exactly one well-founded model, but some ground atoms might be "undefined" in it (they can be defined, however). Thus, the well-founded semantics coincides with the "Consistency Postulate".
- The well-founded model might not be computable (in those not infrequent cases where the fixpoint is reached after more than ω steps).

Example

$S = \{ p(a) \leftarrow \top,\ p(f(x)) \leftarrow p(x),\ q(y) \leftarrow p(y),\ s \leftarrow p(z) \wedge \neg q(z),\ r \leftarrow \neg s \}$

i.e., the (standard) translation of the following set of generalised rules into normal clauses:

$$\{ p(a) \leftarrow \top, \quad p(f(x)) \leftarrow p(x), \quad q(y) \leftarrow p(y), \quad r \leftarrow \forall z (p(z) \Rightarrow q(z)) \}$$

Then

$$lfp(\mathbf{W}_S) = \mathbf{W}_S \uparrow \omega + 2$$
$$= \{ p(a), \quad p(f^n(a)), \quad \} \cup \{ q(a), \quad q(f^n(a)), \quad \} \cup \{ \neg s, r \}$$

Stable and Well-Founded Semantics Compared

- If a rule set is stratifiable, then it has a unique minimal model, which is its only stable model and is also its well-founded model and total.
- If a rule set S has a total well-founded model, then this model is also the single stable model of S.
- If a rule set S has a single stable model, then this model is not necessarily the well-founded model of S.

Example

The set $S = \{p \leftarrow \neg q,\ \ q \leftarrow \neg p,\ \ p \leftarrow \neg p\}$ has the single stable model $\{p\}$, but its well-founded model is $\emptyset$.

- Stable model entailment does *not* imply well-founded entailment:

Example

Let $S = \{p \leftarrow \neg q,\ \ q \leftarrow \neg p,\ \ r \leftarrow p,\ \ r \leftarrow q\}$.
Then r is true in all stable models but it is undefined in the well-founded model.

"reasoning by cases"