

Computational Logic

Datalog with Negation

Free University of Bozen-Bolzano, 2010

Werner Nutt

(Slides by Thomas Eiter and Wolfgang Faber)

The Issue

- Queries like “complement of transitive closure” need both, recursion and negation
- Such queries cannot be expressed in datalog (monotonicity)
- Desired: Extension of datalog with negation

Example: $ready(D) \leftarrow device(D), \neg busy(D)$

- Giving a semantics is not straightforward because of possibly cyclic definitions

Example:

$$\begin{aligned} single(X) &\leftarrow man(X), \neg husband(X) \\ husband(X) &\leftarrow man(X), \neg single(X) \end{aligned}$$

Definition. A *datalog[¬] program* P is a finite set of *datalog[¬] rules* r of the form

$$A \leftarrow B_1, \dots, B_n \tag{1}$$

where $n \geq 0$ and

- A is an atom $R_0(\vec{x}_0)$
- Each B_i is an atom $R_i(\vec{x}_i)$ or a *negated atom* $\neg R_i(\vec{x}_i)$
- $\vec{x}_0, \dots, \vec{x}_n$ are vectors of variables and constants (from **dom**)
- every variable in $\vec{x}_0, \dots, \vec{x}_n$ must occur in some atom $B_i = R_i(\vec{x}_i)$ (“safety”)
- the head of r is A , denoted $H(r)$
- the body of r is $\{B_1, \dots, B_n\}$, denoted $B(r)$, and

$$B^+(r) = \{R(\vec{x}) \mid \exists i B_i = R(\vec{x})\}, B^-(r) = \{R(\vec{x}) \mid \exists i B_i = \neg R(\vec{x})\}$$

P has extensional and intensional relations, $edb(P)$ resp. $idb(P)$, like a datalog program.

Remarks: – “ \neg ” is as in LP often denoted by “not”

– Equality (=) and inequality (\neq , as $\neg =$) are usually available as built-ins, but usage must be “safe”

Datalog[¬] Semantics – The Problem

- **Idea:** Naturally extend the minimal-model semantics of datalog (equivalently, the least fixpoint-semantics) to negation
- Generalize to this aim the immediate consequence operator

$$\mathbf{T}_P(\mathbf{K}): inst(sch(P)) \rightarrow inst(sch(P))$$

Definition. Given a *datalog[¬] program* P and $\mathbf{K} \in inst(sch(P))$,

a fact $R(\vec{t})$ is an *immediate consequence* for \mathbf{K} and P , if either

- $R \in edb(P)$ and $R(\vec{t}) \in \mathbf{K}$, or
- there exists some ground instance r of a rule in P such that
 - * $H(r) = R(\vec{t})$,
 - * $B^+(r) \subseteq \mathbf{K}$, and
 - * $B^-(r) \cap \mathbf{K} = \emptyset$.

(That is, evaluate “ \neg ” w.r.t. \mathbf{K})

Problems with Least Fixpoints

- Natural trial: Define the semantics of datalog[¬] in terms of the least fixpoint of \mathbf{T}_P .

- However, this suffers from several problems:

1. \mathbf{T}_P may not have a fixpoint:

$$P_1 = \{ \textit{known}(a) \leftarrow \neg \textit{known}(a) \}$$

2. \mathbf{T}_P may not have a least (i.e., single minimal) fixpoint:

$$P_2 = \left\{ \begin{array}{l} \textit{single}(X) \leftarrow \textit{man}(X), \neg \textit{husband}(X) \\ \textit{husband}(X) \leftarrow \textit{man}(X), \neg \textit{single}(X) \end{array} \right\}$$

$$\mathbf{I} = \{ \textit{man}(\textit{dilbert}) \}$$

Datalog with Negation

Computational Logic

5

3. The least fixpoint of \mathbf{T}_P including \mathbf{I} may not be constructible by fixpoint iteration (i.e., not as limit $\mathbf{T}_P^\omega(\mathbf{I})$ of $\{\mathbf{T}_P^i(\mathbf{I})\}_{i \geq 0}$):

$$P_3 = P_2 \cup \{ \textit{husband}(X) \leftarrow \neg \textit{husband}(X), \textit{single}(X) \}$$

$$\mathbf{I} = \{ \textit{man}(\textit{dilbert}) \} \text{ as above}$$

Note: the operator \mathbf{T}_P is not monotonic!

Problems with Minimal Models

There are similar problems for model-theoretic semantics

- We can associate with P naturally a first-order theory Σ_P as in the negation-free case (write rules as implications):

$$R(\vec{x}) \leftarrow (\neg)R_1(\vec{x}_1), \dots, (\neg)R_n(\vec{x}_n)$$
$$\rightsquigarrow$$
$$\forall \vec{x} \forall \vec{x}_1 \dots \forall \vec{x}_n ((\neg)R_1(\vec{x}_1) \wedge \dots \wedge (\neg)R_n(\vec{x}_n)) \rightarrow R(\vec{x})$$

- Still, $\mathbf{K} \in inst(sch(P))$ is a model of Σ_P iff $\mathbf{T}_P(\mathbf{K}) \subseteq \mathbf{K}$
(and models are not necessarily fixpoints)
- However, multiple minimal models of Σ_P containing \mathbf{I} might exist
(see the Dilbert example)

Datalog with Negation

Computational Logic

7

Solution Approaches

Different proposals have been made to handle the problems above:

- **Give up single fixpoint/model semantics:** Consider alternative fixpoints (models), and define results by *intersection*, called *certain semantics*.

Most well-known: Stable model semantics (Gelfond & Lifschitz, 1988;1991).

Still suffers from 1.

- **Constrain the syntax of programs:** Consider only a fragment where negation can be “naturally” evaluated to a single minimal model.

Most well-known: semantics for stratified programs (Apt, Blair & Walker, 1988), perfect model semantics (Przymusinski, 1987).

- **Give up 2-valued semantics:** Facts might be true, false or *unknown*

Adapt and refine the notion of immediate consequence.

Most well-known: Well-founded semantics (Ross, van Gelder & Schlipf, 1991).

Resolves all problems 1-3

- **Give up fixpoint/minimality condition:** Operational definition of result.

Most well-known: Inflationary semantics (Abiteboul & Vianu, 1988)

Semi-Positive Datalog

“Easy” case: Datalog \neg programs where negation is applied only to *edb* relations.

- Such programs are called *semi-positive*
- For a semi-positive program, \mathbf{T}_P is monotonic if the *edb*-part is fixed, i.e.,
 $\mathbf{I} \subseteq \mathbf{J}$ and $\mathbf{I}|_{edb(P)} = \mathbf{J}|_{edb(P)}$ implies $\mathbf{T}_P(\mathbf{I}) \subseteq \mathbf{T}_P(\mathbf{J})$

Theorem. Let P be a semi-positive datalog program and $\mathbf{I} \in inst(sch(P))$. Then,

1. \mathbf{T}_P has a unique minimal fixpoint \mathbf{J} such that $\mathbf{I}|_{edb(P)} = \mathbf{J}|_{edb(P)}$.
2. Σ_P has a unique minimal model \mathbf{J} such that $\mathbf{I}|_{edb(P)} = \mathbf{J}|_{edb(P)}$.

Example

Semi-positive datalog can express the transitive closure of the complement of a graph G with vertexes $vert$ and edges $edge$:

$$neg_tc(x, y) \leftarrow vert(x), vert(y), \neg edge(x, y)$$

$$neg_tc(x, y) \leftarrow vert(x), \neg edge(x, z), neg_tc(z, y)$$

Stratified Semantics

- **Intuition:** For evaluating the body of a rule instance r containing $\neg R(\vec{t})$, the value of the “negated” relation $R(\vec{t})$ should be known:

1. evaluate first R
2. if $R(\vec{t})$ is false, then $\neg R(\vec{t})$ is true
3. if $R(\vec{t})$ is true, then $\neg R(\vec{t})$ is false and the rule is not applicable.

- **Example:**

$$boring(chess) \leftarrow \neg interesting(chess)$$
$$interesting(X) \leftarrow difficult(X)$$

For $\mathbf{I} = \{\}$, compute result $\{boring(chess)\}$.

- **Note:** this introduces *procedurality* (violates declarativity)!

Dependency Graph for Datalog⁻ programs

Associate with each datalog⁻ program P a directed graph $DEP(P) = (N, E)$, called *Dependency Graph*, as follows:

- $N = sch(P)$, i.e., the nodes are the relations
- $E = \{\langle R, R' \rangle \mid \exists r \in P: H(r) = R \wedge R' \in B(r)\}$, i.e., arcs $R \rightarrow R'$ from the relations in rule heads to the relations in the body
- Mark each arc $R \rightarrow R'$ with “*”, if $R(\vec{x})$ is in the head of a rule in P whose body contains $\neg R'(\vec{y})$

Remark: *edb* relations are often omitted in the dependency graph


Datalog with Negation

Computational Logic

13

Example

P : $husband(X) \leftarrow man(X), married(X).$
 $single(X) \leftarrow man(X), \neg husband(X).$



Stratification Principle

If $R = R_0 \rightarrow R_1 \rightarrow R_2 \rightarrow \dots \rightarrow R_{n-1} \rightarrow R_n = R'$ such that some $R_i \rightarrow R_{i+1}$ is marked with “*”, then R' must be evaluated prior to R .

Stratification

Definition. A *stratification* of a datalog program P is a partitioning

$$\Sigma = (P_1, \dots, P_n)$$

of $sch(P)$ into nonempty, pairwise disjoint sets P_i such that

1. if $R \in P_i, R' \in P_j$, and $R \rightarrow R'$ is in $DEP(P)$, then $i \geq j$;
2. if $R \in P_i, R' \in P_j$, and $R \rightarrow R'$ is in $DEP(P)$ marked with “*,” then $i > j$.

The sets P_1, \dots, P_n are called the *strata* of P w.r.t. Σ .

Definition. A datalog program P is called *stratified*, if it has some stratification Σ .

Evaluation Order

A stratification Σ defines an *evaluation order* for the relations in P (given $\mathbf{I} \in inst(edb(P))$):

1. First evaluate the relations in P_1 (which is \neg -free).
 \Rightarrow All relations R in heads of P_1 are defined. This yields $\mathbf{J}_1 \in inst(sch(P_1))$.
2. Evaluate P_2 considering relations in $edb(P)$ and P_1 as $edb(P_1)$,
where $\neg R(\vec{t})$ is true if $R(\vec{t})$ is false in $\mathbf{I} \cup \mathbf{J}_1$.
 \Rightarrow All relations R in heads of P_2 are defined. This yields $\mathbf{J}_2 \in inst(sch(P_2))$.
...
3. Evaluate P_i considering relations in $edb(P)$ and P_1, \dots, P_{i-1} as $edb(P_i)$,
where $\neg R(\vec{t})$ is true if $R(\vec{t})$ is false in $\mathbf{I} \cup \mathbf{J}_1 \cup \dots \cup \mathbf{J}_{i-1}$.
4. The result of evaluating P on \mathbf{I} w.r.t. Σ , denoted $P_\Sigma(\mathbf{I})$, is given by $\mathbf{I} \cup \mathbf{J}_1 \cup \dots \cup \mathbf{J}_n$.

$$P = \{ \text{husband}(X) \leftarrow \text{man}(X), \text{married}(X) \\ \text{single}(X) \leftarrow \text{man}(X), \neg \text{husband}(X) \}$$

Stratification Σ :

$$P_1 = \{ \text{man}, \text{married} \}, P_2 = \{ \text{husband} \}, P_3 = \{ \text{single} \}$$

$$\mathbf{I} = \{ \text{man}(\text{dilbert}) \}:$$

1. Evaluate P_1 : $\mathbf{J}_1 = \{ \}$
2. Evaluate P_2 : $\mathbf{J}_2 = \{ \}$
3. Evaluate P_3 : $\mathbf{J}_3 = \{ \text{single}(\text{dilbert}) \}$
4. Hence, $P_\Sigma(\mathbf{I}) = \{ \text{man}(\text{dilbert}), \text{single}(\text{dilbert}) \}$

Datalog with Negation

Computational Logic

17

Formal Definition of Stratified Semantics

Let P be a stratified datalog⁻ program with stratification $\Sigma = (P_1, \dots, P_n)$.

- Let P_i^* be the set of rules from P whose relations in the head are in P_i , and set $edb(P_1^*) = edb(P)$, $edb(P_i^*) = \text{rels}(\bigcup_{j=1}^{i-1} P_j^*) \cup edb(P)$, $i > 1$.
- For every $\mathbf{I} \in \text{inst}(edb(P))$, let $\mathbf{I}_0^\Sigma = \mathbf{I}$ and define

$$\begin{aligned} \mathbf{I}_1^\Sigma &= \mathbf{T}_{P_1^*}^\omega(\mathbf{I}_0^\Sigma) &= \text{lfp}(\mathbf{T}_{P_1^*}(\mathbf{I}_0^\Sigma)) &\supseteq \mathbf{I}_0^\Sigma \\ \mathbf{I}_2^\Sigma &= \mathbf{T}_{P_2^*}^\omega(\mathbf{I}_1^\Sigma) &= \text{lfp}(\mathbf{T}_{P_2^*}(\mathbf{I}_1^\Sigma)) &\supseteq \mathbf{I}_1^\Sigma \\ &\dots && \\ \mathbf{I}_i^\Sigma &= \mathbf{T}_{P_i^*}^\omega(\mathbf{I}_{i-1}^\Sigma) &= \text{lfp}(\mathbf{T}_{P_i^*}(\mathbf{I}_{i-1}^\Sigma)) &\supseteq \mathbf{I}_{i-1}^\Sigma \\ &\dots && \\ \mathbf{I}_n^\Sigma &= \mathbf{T}_{P_n^*}^\omega(\mathbf{I}_{n-1}^\Sigma) &= \text{lfp}(\mathbf{T}_{P_n^*}(\mathbf{I}_{n-1}^\Sigma)) &\supseteq \mathbf{I}_{n-1}^\Sigma \end{aligned}$$

where $\mathbf{T}_Q^\omega(\mathbf{J}) = \lim\{\mathbf{T}_Q^i(\mathbf{J})\}_{i \geq 0}$ with $\mathbf{T}_Q^0(\mathbf{J}) = \mathbf{J}$ and $\mathbf{T}_Q^{i+1} = \mathbf{T}_Q(\mathbf{T}_Q^i(\mathbf{J}))$, and $\text{lfp}(\mathbf{T}_Q(\mathbf{J}))$ is the least fixpoint \mathbf{K} of \mathbf{T}_Q such that $\mathbf{K}|edb(Q) = \mathbf{J}|edb(Q)$.

- Denote $P_\Sigma(\mathbf{I}) = \mathbf{I}_n^\Sigma$

Proposition. For every $i \in \{1, \dots, n\}$,

- $lfp(\mathbf{T}_{P_i^*}(\mathbf{I}_{i-1}^\Sigma))$ exists,
- $lfp(\mathbf{T}_{P_i^*}(\mathbf{I}_{i-1}^\Sigma)) = \mathbf{T}_{P_i^*}^\omega(\mathbf{I}_{i-1}^\Sigma)$,
- $\mathbf{I}_{i-1}^\Sigma \subseteq \mathbf{I}_i^\Sigma$.

Therefore, $P_\Sigma(\mathbf{I})$ is always well-defined.

Stratified semantics singles out a model, and in fact a minimal model.

Theorem. $P_\Sigma(\mathbf{I})$ is a minimal model \mathbf{K} of P such that $\mathbf{K}|_{edb(P)} = \mathbf{I}$.

Dilbert Example (cont'd)

$$P = \{ \text{husband}(X) \leftarrow \text{man}(X), \text{married}(X) \\ \text{single}(X) \leftarrow \text{man}(X), \neg \text{husband}(X) \}$$

$$edb(P) = \{\text{man}\}$$

Stratification Σ : $P_1 = \{\text{man}, \text{married}\}, P_2 = \{\text{husband}\}, P_3 = \{\text{single}\}$

1. $P_1 = \{\}$
2. $P_2 = \{\text{husband}(X) \leftarrow \text{man}(X), \text{married}(X)\}$
3. $P_3 = \{\text{single}(X) \leftarrow \text{man}(X), \neg \text{husband}(X)\}$

$\mathbf{I} = \{\text{man}(\text{dilbert})\}$:

1. $\mathbf{I}_1^\Sigma = \{\text{man}(\text{dilbert})\}$
2. $\mathbf{I}_2^\Sigma = \{\text{man}(\text{dilbert})\}$
3. $\mathbf{I}_3^\Sigma = \{\text{man}(\text{dilbert}), \text{single}(\text{dilbert})\}$

Hence, $P_\Sigma(\mathbf{I}) = \{\text{man}(\text{dilbert}), \text{single}(\text{dilbert})\}$

Stratification Theorem

- The stratification Σ above is not unique.
- Alternative stratification Σ' :
 $P_1 = \{man, married, husband\}, P_2 = \{single\}$
- Evaluation with respect to Σ' yields same result!

The choice of a particular stratification is irrelevant:

Stratification Theorem. Let P be a stratifiable datalog[¬] program. Then, for any stratifications Σ and Σ' and $\mathbf{I} \in inst(sch(P))$, $P_\Sigma(\mathbf{I}) = P_{\Sigma'}(\mathbf{I})$.

- Thus, syntactic stratification yields semantically a canonical way of evaluation.
- The result $P_{str}(\mathbf{I})$ is called the *perfect model* or *stratified model* of P for \mathbf{I} .

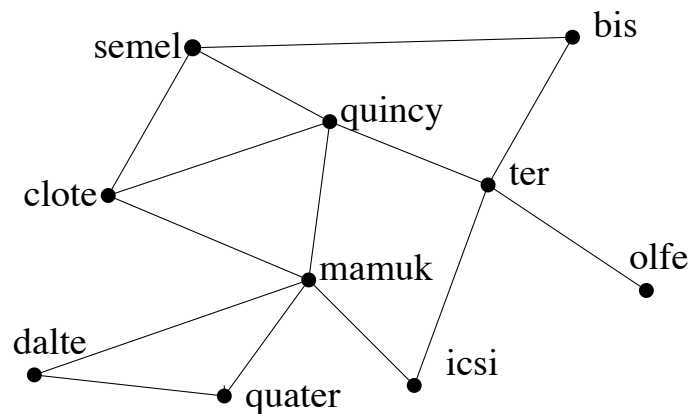
Datalog with Negation

Computational Logic

21

Example: Railroad Network

Determine safe connections between locations in a railroad network



- **Cutpoint c for a and b :** if c fails, there is no connection between a and b
- **Safe connection between a and b :** no cutpoints between a and b exist
- E.g., ter is a cutpoint for $olfe$ and $semel$, while $quincy$ is not.

Datalog with Negation

Relations:

$link(X, Y)$: direct connection from station X to Y (edb facts)

$linked(A, B)$: symmetric closure of $link$

$connected(A, B)$: there is path between A and B (one or more links)

$cutpoint(X, A, B)$: each path from A to B goes through station X

$circumvent(X, A, B)$: there is a path between A and B not passing X

$has_icut_point(A, B)$: there is at least one cutpoint between A and B

$safely_connected(A, B)$: A and B are connected with no cutpoint

$station(X)$: X is a railway station

Railroad program P :

R_1 : $linked(A, B) :- link(A, B)$.

R_2 : $linked(A, B) :- link(B, A)$.

R_3 : $connected(A, B) :- linked(A, B)$.

R_4 : $connected(A, B) :- connected(A, C), linked(C, B)$.

R_5 : $cutpoint(X, A, B) :- connected(A, B), station(X),$
 $\neg circumvent(X, A, B)$.

R_6 : $circumvent(X, A, B) :- linked(A, B), X \neq A, station(X), X \neq B$.

R_7 : $circumvent(X, A, B) :- circumvent(X, A, C), circumvent(X, C, B)$.


R_8 : $has_icut_point(A, B) :- cutpoint(X, A, B), X \neq A, X \neq B$.

R_9 : $safely_connected(A, B) :- connected(A, B),$
 $\neg has_icut_point(A, B)$.

R_{10} : $station(X) :- linked(X, Y)$.

Remark: Inequality (\neq) is used here as built-in. It can be easily defined in stratified manner.

$DEP(P)$:



Stratification Σ :

$$P_1 = \{link, linked, station, circumvent, connected\}$$

$$P_2 = \{cutpoint, has_icut_point\}$$

$$P_3 = \{safely_connected\}$$

Datalog with Negation

Computational Logic

25

$$\mathbf{I}(link) = \{ \langle semel, bis \rangle, \langle bis, ter \rangle, \langle ter, olfe \rangle, \langle ter, icsi \rangle, \langle ter, quincy \rangle, \\ \langle quincy, semel \rangle, \langle quincy, clote \rangle, \langle quincy, mamuk \rangle, \dots, \langle dalte, quater \rangle \}$$

Evaluation $P_\Sigma(\mathbf{I})$:

1. $P_1 = \{link, linked, station, circumvent, connected\}$:

$$\mathbf{J}_1 = linked(semel, bis), linked(bis, ter), linked(ter, olfe), \dots, \\ connected(semel, olfe), \dots, circumvent(quincy, semel, bis), \dots$$

2. $P_2 = \{cutpoint, has_icut_point\}$:

$$\mathbf{J}_2 = cutpoint(ter, semel, olfe), has_icut_point(semel, olfe) \dots$$

3. $P_3 = \{safely_connected\}$:

$$\mathbf{J}_3 = safely_connected(semel, bis), safely_connected(semel, ter)$$

But, $safely_connected(semel, olfe) \notin \mathbf{J}_3$

Algorithm STRATIFY

Input: a datalog[¬] program P .

Output: a stratification Σ for P , or “no” if none exists.

1. construct the directed graph $G := DEP(P) (= \langle N, E \rangle)$ with markers “*”;
2. **for each** pair $R, R' \in N$ **do**
 if R reaches R' via some path containing a marked arc
 then begin $E := E \cup \{R \rightarrow R'\}$; mark $R \rightarrow R'$ with “*” **end**;
3. $i := 1$;
4. identify the set K of all vertices R in N s.t. no marked $R \rightarrow R'$ is in E ;
5. **if** $K = \emptyset$ and G has vertices left, **then** output “no”
 else begin output K as stratum P_i ;
 remove all vertices in K and corresponding arcs from G ;
 end;
6. **if** G has vertices left **then begin** $i := i + 1$; **goto** step 4 **end**
 else stop.

Runs in polynomial time!

Datalog with Negation

Computational Logic

27

Inflationary Semantics for Datalog

Idea: Adopt a production-oriented view of datalog[¬], similar as in rule-based expert systems

- A rule should be applied (fired) if the premises (= body literals) are satisfied with respect to the current state
- Rather than applying one rule at a time (as in expert systems), fire *all* applicable rules in parallel
- New facts may fire other rules
- Repeat application of rules, until no more new facts are generated
- This amounts to the least fixpoint of the inflationary version of $\mathbf{T}_P(\mathbf{K})$

For any datalog[¬] program P , let $\mathbf{T}_P^+ : inst(sch(P)) \rightarrow inst(sch(P))$ denote the inflationary variant of \mathbf{T}_P :

$$\mathbf{T}_P^+(\mathbf{K}) = \mathbf{K} \cup \mathbf{T}_P(\mathbf{K})$$

Definition. Given a datalog[¬] program P and $\mathbf{I} \in inst(edb(P))$, the inflationary semantics of P w.r.t. \mathbf{I} , denoted $P_{inf}(\mathbf{I})$, is the limit of the sequence $\{\mathbf{T}_P^{+i}(\mathbf{I})\}_{i \geq 0}$, where $\mathbf{T}_P^{+0}(\mathbf{I}) = \mathbf{I}$ and $\mathbf{T}_P^{+(i+1)}(\mathbf{I}) = \mathbf{T}_P^+(\mathbf{T}_P^{+i}(\mathbf{I}))$.

Notice:

- $P_{inf}(\mathbf{I})$ is well-defined for each program P and input database \mathbf{I} .
- $P_{inf}(\mathbf{I})$ is a model of P containing \mathbf{I} , but not necessarily a minimal model.
- $P_{inf}(\mathbf{I})$ is a (not necessarily minimal) fixpoint of \mathbf{T}_P^+ containing \mathbf{I} .

Example

$$P = \{q(b) \leftarrow \neg p(a), \quad r(c) \leftarrow \neg q(b) \quad p(a) \leftarrow r(c), \neg p(b)\}$$

Consider $\mathbf{T}_P^{+i}(\mathbf{I})$, $i \geq 0$, for $\mathbf{I} = \emptyset$:

- $\mathbf{T}_P^{+0}(\mathbf{I}) = \mathbf{I} = \{\}$.
- The first two rules are applicable, as $\neg p(a)$, $\neg q(b)$ are satisfied wrt. \mathbf{I}_0 .
- $\mathbf{T}_P^{+1}(\mathbf{I}) = \{q(b), r(c)\}$.
- The third rule is now applicable, as $r(c)$, $\neg p(b)$ are satisfied wrt. \mathbf{I}_1 .
- $\mathbf{T}_P^{+2}(\mathbf{I}) = \{q(b), r(c), p(a)\}$.
- No new facts can be obtained, as all rules have been applied.
- Hence, $P_{inf}(\mathbf{I}) = \mathbf{T}_P^{+2}(\mathbf{I})$.

Note that $P_{inf}(\mathbf{I})$ is not a minimal model of P containing \mathbf{I} .

Example: One-Step-Behind Technique

Undirected graph $G = \langle V, E \rangle$, distance $d: V^2 \longrightarrow \{0, 1, 2, \dots\} \cup \{\infty\}$
($d(x, y)$ = length of shortest path between x, y and ∞ if no path exists)

Define $shorter(x, y, x', y') :\Leftrightarrow dist(x', y') < \infty$

Program P (with $edb(P) = \{v, e\}$, where e is symmetric):

$$t(x, x) \leftarrow v(x)$$

$$t(x, y) \leftarrow t(x, z), e(z, y)$$

$$t1(x, y) \leftarrow t(x, y)$$

$$shorter(x_1, y_1, x_2, y_2) \leftarrow t1(x_1, y_1), t(x_2, y_2), \neg t1(x_2, y_2)$$

$t1(x, y)$ is “one step behind” $t(x, y)$

$$i \geq 0 : \quad t(x, y) \in \mathbf{T}_P^{+i}(\mathbf{I}) \Leftrightarrow dist(x, y) \leq i - 1,$$

$$t1(x, y) \in \mathbf{T}_P^{+i}(\mathbf{I}) \Leftrightarrow dist(x, y) \leq i - 2$$

Datalog with Negation

Computational Logic

31

Inflationary vs Stratified Semantics

- Inflationary Semantics is well-defined for *all* datalog[¬] programs, not only for stratified programs. (It was used e.g. in the FLORID system.)
- For semi-positive programs, inflationary and stratified semantics coincide.
- Datalog[¬] queries under stratified semantics are subsumed by inflationary semantics:

Theorem. For every stratified datalog[¬] program P with “output” relation R , there exists a datalog[¬] program P' such that $edb(P') = edb(P)$ and for all $\mathbf{I} \in inst(edb(P))$,
 $P'_{inf}(\mathbf{I})(R) = P_{strat}(\mathbf{I})(R)$.

- The converse fails, i.e., there are datalog[¬] queries P under inflationary semantics that are not equivalent to any datalog[¬] query under stratified semantics (Kolaitis, 1991).

Intuitive reason: Stratified semantics has a static, fixed number of negation layers, while inflationary semantics allows dynamically many.

Datalog with Negation

- **Idea:** Try to construct a (minimal) fixpoint by iteration from input

If the construction succeeds, the result is the semantics.

- **Problem:** Application of rules might be compromised.

Example:

$$P = \{p(a) \leftarrow \neg p(a), \quad q(b) \leftarrow p(a), \quad p(a) \leftarrow q(b)\}$$

($edb(P)$ is void, thus \mathbf{I} is immaterial and omitted)

- \mathbf{T}_P has the least fixpoint $\{p(a), q(b)\}$
- It is iteratively constructed $\mathbf{T}_P^\omega = \{p(a), q(b)\}$
- $p(a)$ is included into \mathbf{T}_P^1 by the first rule, since $p(a) \notin \mathbf{T}_P^0 = \emptyset$.
- This compromises the rule application, and $p(a)$ is not “foundedly” derived!
- Note: $\mathbf{T}_P^+ = \{p(a), q(b)\}$

Datalog with Negation

Computational Logic

33

Fixed Evaluation of Negation

- **Reason:** \mathbf{T}_P is not monotonic.
- **Solution:** Keep negation throughout fixpoint-iteration fixed.
Evaluate negation w.r.t. a fixed candidate fixpoint model \mathbf{J} .
- Introduce for datalog[¬] program and $\mathbf{J} \in inst(sch(P))$ a new immediate consequence operator $\mathbf{T}_{P,\mathbf{J}}$:

Definition. Given a datalog[⊖] program P and $\mathbf{J}, \mathbf{K} \in inst(sch(P))$, a fact $R(\vec{t})$ is an *immediate* consequence for \mathbf{K} and P under negation \mathbf{J} , if either

- $R \in edb(P)$ and $R(\vec{t}) \in \mathbf{K}$, or
- there exists some ground instance r of a rule in P such that
 - $H(r) = R(\vec{t})$,
 - $B^+(r) \subseteq \mathbf{K}$, and
 - $B^-(r) \cap \mathbf{J} = \emptyset$.

(That is, evaluate “ \neg ” under \mathbf{J} instead of \mathbf{K})

Definition. For any datalog[⊖] program P and $\mathbf{J}, \mathbf{K} \in inst(sch(P))$, let

$$\mathbf{T}_{P,\mathbf{J}}(\mathbf{K}) = \{A \mid A \text{ is an immediate consequence for } \mathbf{K} \text{ and } P \text{ under negation } \mathbf{J}\}$$

Notice:

- $\mathbf{T}_P(\mathbf{K})$ coincides with $\mathbf{T}_{P,\mathbf{K}}(\mathbf{K})$
- $\mathbf{T}_{P,\mathbf{J}}$ is a monotonic operator, hence has for each $\mathbf{K} \in inst(sch(P))$ a least fixpoint containing \mathbf{K} , denoted $lfp(\mathbf{T}_{P,\mathbf{J}}(\mathbf{K}))$
- $lfp(\mathbf{T}_{P,\mathbf{J}}(\mathbf{I}))$ coincides with \mathbf{I} on $edb(P)$ and is the limit $\mathbf{T}_{P,\mathbf{J}}^\omega$ of the sequence

$$\{\mathbf{T}_{P,\mathbf{J}}^i(\mathbf{I})\}_{i \geq 0}$$

where $\mathbf{T}_{P,\mathbf{J}}^0(\mathbf{I}) = \mathbf{I}$ and $\mathbf{T}_{P,\mathbf{J}}^{i+1}(\mathbf{I}) = \mathbf{T}_{P,\mathbf{J}}(\mathbf{T}_{P,\mathbf{J}}^i(\mathbf{I}))$.

Stable Models

Using $\mathbf{T}_{P,\mathbf{J}}$, stable models are defined by requiring that \mathbf{J} is reproduced by the program:

Definition. Let P be a datalog[¬] program P and $\mathbf{I} \in inst(edb(P))$. Then, a stable model for P and \mathbf{I} is any $\mathbf{J} \in inst(sch(P))$ such that

1. $\mathbf{J}|_{edb(P)} = \mathbf{I}$, and
2. $\mathbf{J} = lfp(\mathbf{T}_{P,\mathbf{J}}(\mathbf{I}))$.

Notice: Monotonicity of $\mathbf{T}_{P,\mathbf{J}}$ ensures that at no point in the construction of $lfp(\mathbf{T}_{P,\mathbf{J}}(\mathbf{I}))$ using fixpoint iteration from \mathbf{I} , the application of a rule can be compromised later.

Datalog with Negation

Computational Logic

37

Example

$$P = \{ p(a) \leftarrow \neg p(a), \quad q(b) \leftarrow p(a), \quad p(a) \leftarrow q(b) \}$$

($edb(P)$ is void, thus \mathbf{I} is immaterial and omitted)

- Take $\mathbf{J} = \{p(a), q(b)\}$. Then
 - $\mathbf{T}_{P,\mathbf{J}}^0 = \emptyset$
 - $\mathbf{T}_{P,\mathbf{J}}^1 = \emptyset$
- Thus $lfp(\mathbf{T}_{P,\mathbf{J}}) = \emptyset \neq \mathbf{J}$.
- Hence, the fixpoint \mathbf{J} of \mathbf{T}_P is refuted.
- For P , no stable model exists; thus, it may be regarded as “inconsistent”.

- **Problem:** A datalog program may have multiple stable models:

$$P = \left\{ \begin{array}{l} \text{single}(X) \leftarrow \text{man}(X), \neg \text{husband}(X) \\ \text{husband}(X) \leftarrow \text{man}(X), \neg \text{single}(X) \end{array} \right\}$$

$$\mathbf{I} = \{\text{man}(\text{dilbert})\}$$

- $\mathbf{J}_1 = \{\text{man}(\text{dilbert}), \text{single}(\text{dilbert})\}$ is a stable model:
 - $\mathbf{T}_{P, \mathbf{J}_1}^0(\mathbf{I}) = \{\text{man}(\text{dilbert})\}$
 - $\mathbf{T}_{P, \mathbf{J}_1}^1(\mathbf{I}) = \{\text{man}(\text{dilbert}), \text{single}(\text{dilbert})\}$ (apply 1st rule)
 - $\mathbf{T}_{P, \mathbf{J}_1}^2(\mathbf{I}) = \{\text{man}(\text{dilbert}), \text{single}(\text{dilbert})\} = \mathbf{T}_{P, \mathbf{J}_1}^\omega(\mathbf{I})$
- Similarly, $\mathbf{J}_1 = \{\text{man}(\text{dilbert}), \text{husband}(\text{dilbert})\}$ is a stable model (symmetry)

Datalog with Negation

Computational Logic

39

Stable Model Semantics – Definition

- **Solution:** Define stable semantics of P as the intersection of all stable models (*certain semantics*):

For a datalog[¬] program P and $\mathbf{I} \in \text{inst}(\text{edb}(P))$,

denote by $SM(P, \mathbf{I})$ the set of all stable models for \mathbf{I} and P .

Definition. The stable model semantics of a datalog[¬] program P for $\mathbf{I} \in \text{inst}(\text{edb}(P))$, denoted $P_{sm}(\mathbf{I})$, is given by

$$P_{sm}(\mathbf{I}) = \begin{cases} \bigcap SM(P, \mathbf{I}), & \text{if } SM(P, \mathbf{I}) \neq \emptyset, \\ \mathbf{B}(P, \mathbf{I}), & \text{otherwise.} \end{cases}$$

Examples

•

$$P = \left\{ \begin{array}{l} \text{single}(X) \leftarrow \text{man}(X), \neg \text{husband}(X) \\ \text{husband}(X) \leftarrow \text{man}(X), \neg \text{single}(X) \end{array} \right\}$$

$$P_{sm}(\{\text{man}(\text{dilbert})\}) = \{\text{man}(\text{dilbert})\}$$

•

$$P = \{p(a) \leftarrow \neg p(a), \quad q(b) \leftarrow p(a), \quad p(a) \leftarrow q(b)\}$$

$$P_{sm}(\emptyset) = \{p(a), p(b), q(a), q(b)\} = \mathbf{B}(P, \mathbf{I}).$$

Datalog with Negation

Computational Logic

41

Some Properties

- **Proposition.** Each $\mathbf{K} \in SM(P, \mathbf{I})$ is a minimal model of P such that $\mathbf{K} \upharpoonright_{edb(P)} = \mathbf{I}$.
- **Proposition.** Each $\mathbf{K} \in SM(P, \mathbf{I})$ is a minimal fixpoint of \mathbf{T}_P such that $\mathbf{K} \upharpoonright_{edb(P)} = \mathbf{I}$.
- **Theorem.** If P is a stratified program, then for every $\mathbf{I} \in edb(P)$,
 $P_{sm}(\mathbf{I}) = P_{strat}(\mathbf{I})$.
Thus, stable model semantics extends stratified semantics to a larger class of programs
- Evaluation of stable semantics is intractable: Deciding whether $R(\vec{c}) \in P_{sm}(\mathbf{I})$ for given $R(\vec{c})$ and \mathbf{I} (while P is fixed) is coNP-complete.

Well-Founded Semantics

- **Principle:** Use three truth values: Some facts are true, some false, all others are *unknown*.
- **Intuition:**
 - Positive literals must be derived by applying rules whose body is true
 - Conclude that a negated atom $\neg A$ is true, if A can not be derived by assuming that all facts which are not true are false.

Datalog with Negation

Computational Logic

43

Bibliography

- [1] S. Abiteboul, R. Hull, and V. Vianu. *Foundations of Databases*. Addison-Wesley, 1995.
- [2] K.R. Apt, H.A. Blair, A. Walker, Towards a Theory of Declarative Knowledge, in *Foundations of Deductive Databases and Logic Programming*, J. Minker (ed), pp. 89–148, Morgan Kaufmann, 1988.
- [3] H. Garcia-Molina, J. D. Ullman, and J. Widom. *Database Systems – The Complete Book*. Prentice Hall, 2002.
- [4] DLV homepage, since 1996. <http://www.dbai.tuwien.ac.at/proj/dlv/>.
- [5] M. Gelfond and V. Lifschitz. The Stable Model Semantics for Logic Programming. In *Logic Programming: Proc. Fifth Intl Conference and Symposium*, pp. 1070–1080, 1988. MIT Press.
- [6] M. Gelfond and V. Lifschitz. Classical Negation in Logic Programs and Disjunctive Databases. *New Generation Computing*, 9:365–385, 1991.
- [7] N. Leone, G. Pfeifer, W. Faber, T. Eiter, G. Gottlob, S. Perri, and F. Scarcello. The DLV System for Knowledge Representation and Reasoning. To appear in *ACM Transaction on Computational Logic*. Available at <http://www.arxiv.org/ps/cs.AI/0211004>.
- [8] A. van Gelder, K.A. Ross, J.S. Schlipf, The Well-Founded Semantics for General Logic Programs, *Journal of the ACM*, 38(3):620–650, 1991.

Datalog with Negation