Representing Incomplete Knowledge in Abductive Logic Programming

Marc Denecker Danny De Schreye

Department of Computer Science, K.U.Leuven,
Celestijnenlaan 200A, B-3001 Heverlee, Belgium.
e-mail : {marcd, dannyd}@cs.kuleuven.ac.be

Abstract

Recently, Gelfond and Lifschitz presented a formal language for representing incomplete knowledge on actions and states, and a sound translation from this language to extended logic programming. We present an alternative translation to abductive logic programming with integrity constraints and prove the soundness and completeness. In addition, we show how an abductive procedure can be used, not only for *explanation*, but also for deduction and proving satisfiability under uncertainty.

From a more general perspective, this work can be viewed as a -successful- experiment in the declarative representation of and automated reasoning on incomplete knowledge using abductive logic programming.

1 Introduction

An outstanding problem in logic programming is the representation of incomplete knowledge. In [17], Gelfond and Lifschitz propose an interesting methodology to evaluate the expressivity of a formalism for representing incomplete knowledge. They present a simple temporal language $\mathcal{A}$ for representing incomplete temporal knowledge and apply it successfully to a number of well-known benchmark problems in temporal reasoning. Due to its natural and well-understood semantics, $\mathcal{A}$ provides a direct means to evaluate the expressivity of other formalisms. By providing a transformation from $\mathcal{A}$ to extended logic programming [16] and proving its soundness, they show the expressivity of this formalism for representing temporal knowledge and, more in general, incomplete knowledge.

In the past, another approach has been explored for temporal reasoning, based on event calculus [13], [31], [25], [10]. [10] proposes solutions for the same benchmarks as in [17]. This approach makes use of the formalism of abductive logic programming. One may interpret an abductive program as an *open* logic program in the sense that it contains only definitions for the non-abducible predicates. The completion semantics for abductive logic programs of [3] reflects this well by taking as the completion of an abductive program the set of completed definitions of the non-abducible predicates only. As a consequence, the completion does not impose any restriction on the inter-

pretation of the abducible predicates and incomplete knowledge can be represented via the abducible predicates. Observe that this treatment differs totally from the treatment of a defined predicate with empty definition: the latter predicate is always false.

In general an abductive program will allow too much uncertainty. In most applications, the knowledge engineer has knowledge that implicitly restricts the possible states of the undefined abductive predicates. A declarative way of representing this knowledge is by First-order Logic (FOL) axioms. In order to illustrate this approach, we present a transformation from $\mathcal{A}$ domain descriptions to abductive logic programs with FOL axioms (section 3). Despite the fact that the benchmarks in [10] were solved with abductive event calculus, it turns out to be difficult to transform $\mathcal{A}$ domain descriptions to event calculus programs: as situation calculus, $\mathcal{A}$ is based on the *branching time* philosophy, whereas event calculus is based on a *linear time* philosophy. However, the use of abductive logic programs is in no way restricted to event calculus. The proposed transformation maps an $\mathcal{A}$ domain description to an abductive situation calculus.

A second high level goal is to illustrate how an abductive procedure can be useful for automated reasoning with abductive logic programs and FOL axioms. That an abductive procedure can be used for *explanation* of some observation is well-known from [29], [31]. It is less known that an abductive procedure can also be used deduction and for proving consistency of a theory. In section 4, we illustrate how an abductive procedure can be used to prove the consistency of an abductive logic program with FOL axioms and that an abductive procedure can be used for deduction in an abductive logic program.

In the rest of the paper we will mostly refer to *abductive logic programs* as *open logic programs*, and to *abducible predicates* as *undefined predicates*. We prefer these names over the commonly used terminology because they reflect the declarative meaning of the concepts: an undefined predicate has no definition; an open logic program is open in the sense that it describes the defined predicates in terms of a set of undefined predicates whose interpretation is completely left open. The terminology *abductive logic program* and *abducible predicate* is somewhat misleading, especially in the context of this paper, whose goal it is to show explicitly how to perform other forms of reasoning than abduction on *abductive logic programs*.

The paper is structured as follows. In section 2, we recall the language $\mathcal{A}$ and its semantics. In section 3, the transformation from $\mathcal{A}$ to situation calculus programs is presented and the soundness and completeness are proved. In section 4, the use of abduction for explanation and deduction is illustrated. Section 5 gives a comparison between our transformation and the transformation in [17]. Section 6 compares our work with another recent transformation from $\mathcal{A}$ proposed in [11]. In section 7, we discuss other related work. A short paper on this subject has been published as [8].

2 The temporal language $\mathcal{A}$

The language $\mathcal{A}$ [17] allows to describe relationships between fluents (= time dependent properties of the world) and actions. $\mathcal{A}$ is a propositional language: both fluents and actions are represented by propositional symbols. Two types of expressions occur.

A v-proposition describes the value of a fluent after a (possibly empty) sequence of actions. Its syntax is as follows:

f after $a_1; \dots; a_n$

Here $a_1; \dots; a_n$ is a sequence of action symbols and f is a fluent expression: a positive or negative literal containing a fluent. The expression means that f is true after executing the sequence of actions $a_1; \dots; a_n$. If the sequence of actions is empty ($n=0$), the v-proposition describes the initial situation. Instead of **f after**, one usually writes:

Initially f

An e-proposition describes the effect of actions on the fluents. It has the form:

a causes f if $p_1, \dots, p_n$

where $f, p_1, \dots, p_n$ are fluent expressions and a is an action symbol. The expression means that if $p_1, \dots, p_n$ are true, then the effect of a on the current situation is that f becomes true. $p_1, \dots, p_n$ are called preconditions. If $n=0$, one writes:

a causes f .

A domain description is a set of v- and e-propositions.

Example We recall the Yale Turkey Shooting problem (YTS) as formulated in [17].

The fluents are *loaded*, *alive*; the action names are *shoot*, *wait* and *load*. The domain description D_{YTS} contains the following propositions:

Initially *alive*
Initially $\neg$ *loaded*
load* causes *loaded
shoot* causes $\neg$ *alive* if *loaded
shoot* causes $\neg$ *loaded

Example The Murder Mystery domain D_{MM} is a variant of YTS, obtained by substituting

$\neg$ *alive* after *shoot*; *wait*

for **Initially $\neg$ *loaded*** in D_{YTS} . This is a prototypical postdiction problem: we want to obtain the conclusion that initially the gun is loaded.

The semantics for $\mathcal{A}$ is defined as follows. A *state* is a set of fluent names and describes a possible state of the world. Given a fluent symbol f and a state σ , f holds in σ if $f \in \sigma$, otherwise $\neg f$ holds in σ . A transition function Φ maps pairs (a, σ) of action symbols a and states σ into the set of states. Φ describes how a situation changes under application of an action a . A structure M is a pair (σ_0, Φ) , where σ_0 represents the initial state and Φ the transition function. $M^{a_1; \dots; a_m}$ denotes the state $\Phi(a_m, \Phi(a_{m-1}, \dots, \Phi(a_1, \sigma_0) \dots))$. A v-proposition

f after $a_1; \dots; a_m$

holds in a structure (σ_0, Φ) iff f holds in $M^{a_1; \dots; a_m}$.

Definition 2.1 A structure $M = (\sigma_0, \Phi)$ is a model of a domain expression D if and only if the following rules are satisfied:

- Each v -proposition f after $a_1; \dots; a_m$ holds in M .
- For any state σ , fluent expression f and action a , if there exists an e -proposition

a causes f if $p_1, \dots, p_n$

such that $p_1, \dots, p_n$ hold in σ , then f holds in $\Phi(a, \sigma)$. Otherwise, f holds in $\Phi(a, \sigma)$ iff f holds in σ .

A domain description is called *consistent* if it has a model. We introduce a new notion, *e-consistency*. A domain description is *e-consistent* if the set of e -propositions of D is consistent. There is a simple necessary and sufficient condition for a domain description to be e -consistent.

Lemma 2.1 A domain description D is e -consistent iff for each pair of rules

a causes f if $p_1, \dots, p_n$ and a causes $\neg f$ if $p_{n+1}, \dots, p_m$

in D , there exists an i and j such that p_i is the complement of p_j .

This condition is satisfied when the complementary literals are found in the bodies of the two rules but also when they appear in the body of one rule, as in

shoot causes alive if loaded, $\neg$ loaded

Such a rule has an inconsistent body. It can never be applied and can never cause an inconsistency.

In [11], an action a for which there exists a pair of rules

a causes f if $p_1, \dots, p_n$ and a causes $\neg f$ if $p'_1, \dots, p'_m$

such that no complement of a literal in the first is contained in the second is called self-contradicted. [11] contains a similar proposition as lemma 2.1.

Proof Assume that for some pair of rules:

a causes f if $p_1, \dots, p_n$ and a causes $\neg f$ if $p_{n+1}, \dots, p_m$

no complementary literals occur in $\{p_1, \dots, p_m\}$. This set may contain positive and negative literals. Define the state σ consisting of only the positive literals in $\{p_1, \dots, p_m\}$. Then obviously each p_i holds in σ . $\Phi(a, \sigma)$ is not consistently defined because both f and $\neg f$ should hold in $\Phi(a, \sigma)$.

Vice versa, one easily verifies that when the syntactical condition in the lemma is satisfied, the definition of a model of a domain description gives a consistent description of a transition function Φ . $\square$

Definition 2.2 (Non-inertial) We say that a fluent f is non-inertial under an action a in a state σ iff there exists an e-proposition

$$a \text{ causes } f \text{ if } p_1, \dots, p_n \text{ or } a \text{ causes } \neg f \text{ if } p_1, \dots, p_n$$

such that $p_1, \dots, p_n$ hold in σ . Otherwise, it is inertial under action a .

[17] observes that the e-propositions of a domain D completely determine the transition function Φ . Below proposition 2.1 gives a precise and deterministic characterisation of Φ , under the condition that D is e-consistent.

Proposition 2.1 Let D be e-consistent. If D has a model (σ_0, Φ) then Φ is defined as follows: for any state σ , action a and fluent symbol f : $f \in \Phi(a, \sigma)$ iff

- f holds in σ and f is inertial under a in state σ , or
- there exists an e-proposition $a \text{ causes } f \text{ if } p_1, \dots, p_n$ such that $p_1, \dots, p_n$ hold in σ .

The proof is straightforward.

Note that if the condition of e-consistency is not satisfied, then the description of Φ in the proposition does not correspond to the description of Φ in definition 2.1. Indeed, for f , a and σ such that for e-propositions

$$a \text{ causes } f \text{ if } p_1, \dots, p_n \text{ and } a \text{ causes } \neg f \text{ if } p'_1, \dots, p'_m$$

the fluents $p_1, \dots, p_n, p'_1, \dots, p'_m$ hold in σ , the second item of definition 2.1 requires that both $f \in \Phi(a, \sigma)$ and $f \notin \Phi(a, \sigma)$. This is a contradiction. On the other hand, proposition 2.1 requires that $f \in \Phi(a, \sigma)$ since the second item applies. Because the transformation proposed in next section implements the formulation of proposition 2.1 rather than that of definition 2.1, it will be complete only for e-consistent domains.

A v-proposition Q is *entailed* by a domain description D iff Q holds in each model of D . A domain description is called *complete* if it has a unique model. The YTS domain and Murder Mystery domain are examples of complete domain descriptions. Since they share their e-propositions, their models have an identical transition function Φ which maps tuples $(wait, \sigma)$ on σ , $(load, \sigma)$ on $\sigma \cup \{loaded\}$, and maps $(shoot, \sigma)$ on $\sigma \setminus \{alive, loaded\}$ if $loaded \in \sigma$, otherwise on σ . The model M_0 of D_{YTS} has initial situation $\{alive\}$. The model M_1 of D_{MM} has initial situation $\{alive, loaded\}$. An *incomplete* domain description is obtained by dropping the v-proposition **Initially alive** from the Murder Mystery domain. One additional model is the structure with transition function Φ but with initial situation $\{\}$.

$\mathcal{A}$ provides only restricted expressivity: the language is only propositional, no relationships between fluents can be defined, no indeterminate events are allowed. Nevertheless, $\mathcal{A}$ allows to formalise several interesting domains. This and its clear semantics makes the language interesting for experiments as in [17], [11] and in this paper.

3 Translation to Open Logic Programs

In this section we present a general translation from an $\mathcal{A}$ domain description D to an open logic program with FOL axioms. An open logic program is a pair of a set of normal clauses, Horn clauses augmented with *negation as failure* and a set of predicates, called undefined or abducible, such that the undefined predicates occur only in the body of the rules. Intuitively, an open logic program allows to represent incomplete information because it leaves open the interpretation of the undefined predicates. To see this, consider the completion semantics of Console, Theseider Dupré and Torasso for abductive logic programs [3], an extension of Clark's completion semantics. According to this completion semantics, the declarative meaning of an open logic program P is given by the theory, denoted $comp_a(P)$, which consists of the axioms of Free Equality (FEQ , also called Clark Equality Theory or Unique names axioms) and the completed definitions of all *defined* predicates, i.e. all predicates of the language which are not undefined. Otherwise said, $comp_a(P)$ is obtained by taking P 's Clark completion and dropping for each undefined predicate p/n its completed definition $p(\overline{X}) \leftrightarrow false$.

Note that $comp_a(P)$ contains only equivalences with defined predicates at the left. Intuitively, such an equivalence defines when the predicate at the left is true in terms of the situations described at the right. Since undefined predicates have no such *defining* equivalence, the logic program leaves their interpretation open to a great extend¹. In general, an abductive logic program allows too much freedom to the interpretation of undefined predicates. Most often, the program P must be augmented with a set T of First-order Logic (FOL) axioms representing other information which restricts the state of the undefined predicates. The declarative semantics of a couple (P, T) is given by the FOL theory $comp_a(P) + T$. This defines indirectly a model semantics for (P, T) : in the sequel, when we talk about a model of (P, T) we mean a classical model of $comp_a(P) + T$.

That the resulting formalism is adequate for representing incomplete information is now obvious: for $P = \{\}$ with all predicates undefined, the theory $comp_a(P) + T$ collapses to $FEQ + T$, i.e. classical logic with unique names axioms. The expressivity of this formalism for representing incomplete information is widely accepted.

Our transformation from $\mathcal{A}$ to the above formalism produces programs in situation calculus style. Traditionally, two options are available to represent a fluent f in a logic formalism: by a predicate $f(s)$ or by $Holds(f, s)$ where s is a state argument. Then $\neg f$ is translated to $\neg f(s)$ or $\neg Holds(f, s)$. The two approaches are equivalent but the meta-approach has the advantage that the frame axiom can be stated for all fluents at once, whereas in the first approach one frame axiom per fluent predicate is needed. As

¹ As so often with completion semantics, examples can be formulated which contradict this intuition. For example, the 2-valued completion of the program:

$$p \text{ :- } \neg p, \neg a$$

with undefined predicate a entails that a is true. The problem is caused by the loop over negation of p . For programs which do not contain such loops, it can be shown that the undefined predicates can have any interpretation. In other semantics such as the 3-valued completion semantics for abductive programs [4], the justification semantics for abductive programs [7] and the generalised well-founded semantics for abductive logic programs [28], even for programs with loops over negation, the interpretation of the undefined predicates can be any. Despite these problems with 2-valued completion semantics, we use it here because of its declarative simplicity and its close relationship with First-order Logic.

in [17], we use *Holds*/2. The first order language $\mathcal{L}_D$ contains the predicate symbols *Holds*/2, *Noninertial*/3 and *Initially*/1. Each fluent and action symbol occurs in $\mathcal{L}_D$ as a constant. In addition, there is a constant s_0 to denote the initial state and a functor *Result*/2: *Result*(a, s) denotes the state obtained by applying action a on state s . In the sequel, we will use *Result*[$a_1; \dots; a_n, s$] as a shorthand notation for *Result*($a_n, \text{Result}(a_{n-1}, \dots, \text{Result}(a_1, s) \dots)$). For $n = 0$, this denotes s .

$\mathcal{A}$ allows uncertainty on the initial state. Correspondingly, the open program comprises one undefined predicate, *Initially*/1. The translation maps a domain description D to a theory πD consisting of an open logic program P_D and a set of FOL axioms IC_D . P_D is defined as follows:

- Initialisation:

$$(3.1) \quad \text{Holds}(F, s_0) :- \text{Initially}(F)$$

- Law of Inertia:

$$(3.2) \quad \text{Holds}(F, \text{Result}(A, S)) :- \text{Holds}(F, S), \neg \text{Noninertial}(F, A, S)$$

- For each e-proposition a **causes** f **if** $p_1, \dots, p_m, \neg p'_1, \dots, \neg p'_n$ with f, p_i and p'_j positive literals:

$$(3.3) \quad \text{Holds}(f, \text{Result}(a, S)) :- \text{Holds}(p_1, S), \dots, \text{Holds}(p_m, S), \\ \neg \text{Holds}(p'_1, S), \dots, \neg \text{Holds}(p'_n, S)$$

As in [17], we introduce the convention that when f is a negative literal $\neg f'$, *Holds*(f, t) is used as a textual denotation for $\neg \text{Holds}(f', t)$. This handsome convention allows us to say that a **causes** f **if** $p_1, \dots, p_n$ is translated to the clause:

$$\text{Holds}(f, \text{Result}(a, S)) :- \text{Holds}(p_1, S), \dots, \text{Holds}(p_n, S)$$

without considering the sign of the literals p_i . Be aware that a program should never contain literals of the form *Holds*($\neg f, S$), and that when these literals are found in this and the following section, they always stand for $\neg \text{Holds}(f, S)$.

- For each e-proposition a **causes** f **if** $p_1, \dots, p_n$ with f a positive or negative fluent literal:

$$(3.4) \quad \text{Noninertial}(|f|, a, S) :- \text{Holds}(p_1, S), \dots, \text{Holds}(p_n, S)$$

where $|\cdot|$ maps a fluent expression to the comprised fluent symbol, i.e. for a fluent symbol f , both $|f|$ and $|\neg f|$ denote f .

The set of FOL axioms IC_D is defined as follows:

- For each v-proposition f **after** $a_1; \dots; a_n$ ($n \geq 0$):

$$(3.5) \quad \text{Holds}(f, \text{Result}[a_1; \dots; a_n, s_0])$$

with the same syntactic convention on $Holds/2$ as above.

Example The domain description D_{YTS} for the YTS problem is transformed to:

$$\begin{aligned} &Holds(F, s_0) :- Initially(F) \\ &Holds(F, Result(A, S)) :- Holds(F, S), \neg Noninertial(F, A, S) \\ &Holds(loaded, Result(load, S)) :- \\ &Noninertial(loaded, load, S) :- \\ &Noninertial(loaded, shoot, S) :- \\ &Noninertial(alive, shoot, S) :- Holds(loaded, S) \\ \\ &Holds(alive, s_0) \\ &\neg Holds(loaded, s_0) \end{aligned}$$

The clause $Noninertial(loaded, load, S) :-$ may be dropped from this program without effect on the semantics of $Holds/2$. In general, all $Noninertial/3$ rules for initiating effects of actions may be dropped, without effect on the semantics of $Holds/2$.

πD_{YTS} strongly resembles the YTS solution in [1] and in [14]. For example, [1] proposes a Prolog program analogous to πD_{YTS} :

$$\begin{aligned} &Holds(alive, s_0) :- \\ &Holds(F, Result(A, S)) :- Holds(F, S), \neg Noninertial(F, A, S) \\ &Holds(loaded, Result(load, S)) :- \\ &Noninertial(loaded, shoot, S) :- \\ &Noninertial(alive, shoot, S) :- Holds(loaded, S) \end{aligned}$$

Note that this program entails the two FOL axioms. [1] proves that the program is acyclic. The same holds for πD_{YTS} , and in fact for all transformed domain descriptions:

Proposition 3.1 *The translation πD of any domain description D is acyclic.*

Proof A slight modification of the level mapping proposed in [1] for the YTS program applies for all domain descriptions. For all ground terms t , let $|t|_{Result}$ denote the number of occurrences of the functor $Result/2$ in t . We define $|\cdot|$ for all ground terms t, a and s as follows:

$$\begin{aligned} |Initially(t)| &= 0 \\ |Holds(t, s)| &= 2 \times |s|_{Result} + 1 \\ |Noninertial(t, a, s)| &= 2 \times |s|_{Result} + 2 \end{aligned}$$

One easily verifies that $|\cdot|$ is a level mapping. □

Several types of semantics have been defined for open logic programs: 2-valued completion semantics [3], generalised stable semantics [20], the generalised well-founded semantics [28], 3-valued completion semantics and 3-valued (direct) (partial) justification semantics with FEQ [7]. Due to the fact that πD is acyclic and in each clause of P_D , the variables of the body occur in the head, all these semantics coincide in the (weak) sense that the set of all ground atoms implied by πD under any of the semantics is identical. This extension of results of [1] is proven formally in [4]. Because of

these results, we can investigate the soundness and completeness of the transformation under the simplest semantics for open logic programs, the completion semantics of [3].

The translation πD of a domain description contains two defined predicates, $Holds/2$ and $Noninertial/3$. The completed definition of $Holds/2$ is of the form:

$$(3.6) \quad \forall F, T : Holds(F, T) \leftrightarrow E_1 \vee E_2 \vee \dots \vee E_m$$

$$\begin{aligned} \text{with: } E_1 &= T = s_0 \wedge Initially(F) \\ E_2 &= \exists A, S : T = Result(A, S) \wedge \\ &\quad Holds(F, S) \wedge \neg Noninertial(F, A, S) \\ E_i (i > 2) &= \exists S : F = f \wedge T = Result(a, S) \wedge \\ &\quad Holds(p_1, S) \wedge \dots \wedge Holds(p_n, S) \end{aligned}$$

such that there is a correspondence between the disjuncts E_i ($i > 2$) and the e-propositions $a \text{ causes } f \text{ if } p_1, \dots, p_n$ with f a fluent symbol.

If we map S on state σ , the *Result* functor on Φ and $Holds(F, S)$ to $F \in \sigma$ then the completed definition of $Holds/2$ is similar to proposition 2.1.

The completed definition of $Noninertial/3$ is of the form:

$$(3.7) \quad \forall A, F, S : Noninertial(F, A, S) \leftrightarrow E_1 \vee \dots \vee E_m$$

$$\text{with: } E_i = A = a \wedge F = |f| \wedge Holds(p_1, S) \wedge \dots \wedge Holds(p_n, S)$$

such that precisely for each e-proposition $a \text{ causes } f \text{ if } p_1, \dots, p_n$ (f positive or negative), there exists one corresponding disjunct in the completed definition. This formula is the counterpart of definition 2.2: the formulation is almost identical apart from the fact that $Holds(p_i, S)$ should be replaced by " p_i holds in S ".

A transformation such as π can be considered correct if the set of entailed formulas are equivalent. In [17] a translation π is defined to be sound iff for each domain D and v-expression Q , if $\pi D \models \pi Q$ then D entails Q . π is defined to be complete if the reverse holds: if D entails Q then $\pi D \models \pi Q$.

Theorem 3.1 (soundness) *Let D be a domain description. For any v-proposition Q , if $\pi D \models \pi Q$ then D entails Q .*

For the proof of the theorem, we introduce the concept of a state function.

Definition 3.1 *Let M be any (possibly non-Herbrand) interpretation of $\mathcal{L}_D$, x an element of the domain of M .*

$$state_M(x) = \{f \mid f \text{ is a fluent symbol and } M \models Holds(f, x)\}$$

The notation in which domain elements of a model appear in a term or formula, as in $Result(a, x)$ and $Holds(f, x)$, is not conventional. Such a term or formula can be interpreted as a pair of a term or a formula with free variables and a variable assignment of these free variables. E.g. $Holds(f, x)$ corresponds to the pair $(Holds(f, X), \{X/x\})$. " $M \models Holds(f, x)$ " corresponds to the more conventional " $M, \{X/x\} \models Holds(f, X)$ ".

Note that $state_M(x)$ denotes a state as in section 2, and a transition function Φ can be applied on it. Note also that $f \in state_M(x)$ is equivalent with $M \models Holds(f, x)$.

Proof of Theorem 3.1.

Nothing is to be proved when D is inconsistent. So, assume D is consistent. It suffices to show that for each model $M_a = (\sigma_0, \Phi)$ of D and for each v-proposition Q , there exists a model M of πD such that Q holds in (σ_0, Φ) iff $M \models \pi Q$. If then πQ holds in all models of πD , then also in these models corresponding to models of D . Hence, Q holds in all models of D .

We construct a Herbrand model M . $HU(\mathcal{L}_D)$ denotes the Herbrand universe. The basic idea is simple: we define M such that for each state term $s = \text{Result}[a_1; \dots; a_n, s_0]$, $\text{state}_M(s) = M^{a_1; \dots; a_n}$ and such that the atom $\text{Noninertial}(a, f, s)$ is true whenever f is non-inertial under a in the state $\text{state}_M(s)$. Things are complicated due to the fact the πD is not a sorted program and ill-sorted atoms may occur in the model. For that reason, we extend D to D' by allowing each term $t \in HU(\mathcal{L}_D)$ as a fluent symbol and as an action symbol. Note that non-original fluents t_f (symbols of $HU(\mathcal{L}_D)$ which are not fluent symbols in D) and non-original actions t_a do not occur in the e-propositions. This implies that a non-original fluent t_f of D' always remains as in the initial state; a non-original action t_a has no effect on a state (like the action *wait* in the YTS D_{YTS}). (σ_0, Φ) can easily be extended to a model (σ'_0, Φ') of D' . Define $\sigma'_0 = \sigma_0$. Let σ' be any state of D' , consisting of original fluents σ and new fluents σ_n . Extend Φ to Φ' in the following way:

$$\begin{aligned}\Phi'(t_a, \sigma') &= \sigma' && \text{for } t_a \text{ a non-original action} \\ \Phi'(t_a, \sigma') &= \Phi(t_a, \sigma) \cup \sigma_n && \text{for } t_a \text{ an original action}\end{aligned}$$

One easily verifies that (σ'_0, Φ') is a model of D' . Moreover, for any v-proposition Q based on the original language of D , Q holds in (σ'_0, Φ') iff Q holds in (σ_0, Φ) .

Next we associate to (σ'_0, Φ') a Herbrand model M of $\pi D = \pi D'$. With any term $t_s \in HU(\mathcal{L}_D)$ we associate a specific state of D' . Below we call a term t_s an empty-state term if $t_s \neq s_0$ and $t_s \neq \text{Result}(t_1, t_2)$ for some t_1, t_2 . For any term $t_s \in HU(\mathcal{L}_D)$, $\text{state}_M(t_s)$ is constructed as follows:

$$\begin{aligned}t_s = s_0 &\Rightarrow \text{state}_M(t_s) = \sigma_0 \\ t_s \text{ is an empty-state term} &\Rightarrow \text{state}_M(t_s) = \phi \\ t_s = \text{Result}[t_1; \dots; t_n, s_0] &\Rightarrow \text{state}_M(t_s) = \Phi'[t_1; \dots; t_n, \sigma_0] \\ t_s = \text{Result}[t_1; \dots; t_n, t_0], t_0 \text{ an empty-state term} &\Rightarrow \text{state}_M(t_s) = \Phi'[t_1; \dots; t_n, \phi]\end{aligned}$$

M is defined as follows:

$$\begin{aligned}&\{\text{Holds}(t_f, t_s) \mid t_f \in \text{state}_M(t_s)\} \cup \\&\{\text{Noninertial}(t_f, t_a, t_s) \mid t_f \text{ is non-inertial under } t_a \\&\quad \text{in } \text{state}_M(t_s)\} \cup \\&\{\text{Initially}(t_f) \mid t_f \in \sigma_0\}\end{aligned}$$

Clearly for any v-proposition Q using original symbols of D , it holds that Q holds in (σ_0, Φ) iff $M \models \pi Q$. A direct consequence is that M is a model of IC_D . It remains to be proven that M is a model of $\text{comp}_a(P_D)$.

Before continuing with this proof, we want to stress that the complexity of the construction above is in no way an indication that the proposed transformation π is on itself unnecessarily complex or lacks elegance. The increased technicality is only due to the fact that πD can be considered as an untyped meta-program. It is well-known (see e.g. [19], [24]) that such programs give rise to technical problems with respect to Herbrand semantics. Alternatives would have been to define πD as a typed logic program (as in [11]), or to make its clauses range restricted, using additional range predicates. Both solutions would have reduced the complexity of the proof, but increased the complexity of π itself. This motivates our choice.

That the completed definition (3.7) of *Noninertial*/3 is satisfied follows straightforwardly: since the expression " p_i holds in $state_M(t_s)$ " is equivalent with $M \models Holds(p_i, t_s)$, the completed definition of *Noninertial*/3 is a direct representation of definition 2.2.

Finally, we check the completed definition (3.6) of *Holds*/2. Essentially what must be done is to check all its ground instances $Holds(t_f, t_s) \leftrightarrow \dots$. This requires a simple case-analysis depending on the type of t_s . We consider three cases. Take $t_s = s_0$. The completed definition collapses to:

$$Holds(t_f, s_0) \leftrightarrow Initially(t_f)$$

which is clearly satisfied in M .

Take t_s an empty-state term. The completed definition collapses to:

$$Holds(t_f, t_s) \leftrightarrow false$$

which is also satisfied in M .

Finally, take $t_s = Result(t_a, t)$. The completed definition collapses to the equivalence Eq1:

$$Holds(t_f, Result(t_a, t)) \leftrightarrow E_1 \vee \dots \vee E_n$$

$$\begin{aligned} \text{with: } E_1 &= Holds(t_f, t) \wedge \neg Noninertial(t_f, t_a, t) \\ E_i (i > 1) &= Holds(p_1, t) \wedge \dots \wedge Holds(p_n, t) \end{aligned}$$

such that there is a correspondence between the disjuncts $E_i (i > 1)$ and the e-propositions t_a **causes** t_f **if** $p_1, \dots, p_n$ in D .

Now, recall that (σ_0, Φ') is a model of D' . By proposition 2.1, the equivalence Eq2 holds:

$$\begin{aligned} t_f \in \Phi'(t_a, state_M(t)) \text{ iff} \\ \bullet \ t_f \text{ holds in } state_M(t) \text{ and } t_f \text{ is inertial under } t_a \text{ in } state_M(t), \text{ or} \\ \bullet \text{ there exists an e-proposition } t_a \text{ causes } t_f \text{ if } p_1, \dots, p_n \text{ such that} \\ \quad p_1, \dots, p_n \text{ hold in } state_M(t). \end{aligned}$$

Now compare the equivalences Eq1 and Eq2. Because of the following equivalences:

$$\begin{aligned}
M &\models \text{Holds}(p_i, t) \text{ iff } p_i \text{ holds in } \text{state}_M(t) \\
M &\models \neg \text{Noninertial}(t_f, t_a, t) \text{ iff } t_f \text{ is inertial under } t_a \text{ in } \text{state}_M(t)
\end{aligned}$$

one easily verifies that the right-hand of Eq2 is satisfied iff the right-hand of Eq1 is satisfied in M .

The left-hand of Eq1 holds in M iff:

$$M \models \text{Holds}(t_f, \text{Result}(t_a, t)) \text{ iff } t_f \in \text{state}_M(\text{Result}(t_a, t))$$

Now observe that by the definition of M , we have

$$t_f \in \text{state}_M(\text{Result}(t_a, t)) \text{ iff } t_f \in \Phi'(t_a, \text{state}_M(t))$$

Hence, the left-hand of Eq1 is satisfied in M iff the left-hand of Eq2 is true. Because Eq2 is true, Eq1 is satisfied in M .

□

Theorem 3.2 (completeness) *Let D be e-consistent. For each v-proposition Q , if D entails Q then $\pi D \models \pi Q$.*

Proof Since D is e-consistent, there exists a unique transition function Φ which satisfies the e-propositions of D . As for the soundness, it suffices to prove that for each model M of πD , there exists a model $M_a = (\sigma_0, \Phi)$ of D such that for each v-proposition Q , $M \models \pi Q$ iff Q holds in M_a . Notice that this immediately implies that all v-propositions of D hold in M_a since M is a model of πQ for each v-proposition Q of D .

M maps each term $\text{Result}[a_1; \dots; a_n, s_0]$ to a domain element which we denote as:

$$\tilde{M}(\text{Result}[a_1; \dots; a_n, s_0])$$

M_a is defined in the following way: Φ is given; σ_0 is defined as $\text{state}_M(\tilde{M}(s_0))$.

We should prove that for each sequence of actions $a_1, \dots, a_n$:

$$M_a^{a_1; \dots; a_n} = \text{state}_M(\tilde{M}(\text{Result}[a_1; \dots; a_n, s_0]))$$

The proof is by induction on n . For $n = 0$, this is trivial. So assume that the theorem holds for $n - 1$, $n > 0$. We have the following identity:

$$\begin{aligned}
M_a^{a_1; \dots; a_n} &= \Phi(a_n, M_a^{a_1; \dots; a_{n-1}}) \\
&= \Phi(a_n, \text{state}_M(\tilde{M}(\text{Result}[a_1; \dots; a_{n-1}, s_0])))
\end{aligned}$$

The second identity follows from the induction hypothesis. Let x be the domain element $\tilde{M}(\text{Result}[a_1; \dots; a_{n-1}, s_0])$. It suffices to show that:

$$\begin{aligned}
\Phi(a_n, \text{state}_M(\tilde{M}(\text{Result}[a_1; \dots; a_{n-1}, s_0]))) &= \\
&\text{state}_M((\tilde{M}(\text{Result}[a_1; \dots; a_n, s_0])))
\end{aligned}$$

or equivalently:

$$\Phi(a_n, state_M(x)) = state_M(\tilde{M}(Result(a_n, x)))$$

By proposition 2.1, we find that $f \in \Phi(a_n, state_M(x))$ iff

- f holds in $state_M(x)$ and f is inertial under a_n in $state_M(x)$, or
- there exists an e-proposition a_n **causes** f if $p_1, \dots, p_m$ such that the fluents $p_1, \dots, p_m$ hold in $state_M(x)$.

Because M is a model of *Noninertial/3*, the first disjunct corresponds to

$$M \models Holds(f, x) \wedge \neg Noninertial(f, a_n, x)$$

The second disjunct corresponds to the fact that

$$M \models Holds(p_1, x) \wedge \dots \wedge Holds(p_m, x)$$

for some e-proposition a_n **causes** f if $p_1, \dots, p_m$. Because M is a model of the completed definition of *Holds/2*, we obtain that

$$f \in \Phi(a_n, state_M(x)) \text{ iff } M \models Holds(f, Result(a_n, x))$$

or equivalently

$$f \in state_M(\tilde{M}(Result(a_n, x)))$$

This gives the desired identity. $\square$

The following example shows that the condition of e-consistency is necessary: π is not complete in general.

Example Consider the following domain description D_2 , which uses the fluent *alive* and the action *shoot*.

shoot **causes** *alive*
shoot **causes** $\neg$ *alive*

Obviously, D_2 is inconsistent: no transition function Φ can exist which satisfies the two e-propositions. Therefore, each v-proposition is entailed by D_2 . πD_2 is given by:

$Holds(F, s_0) :- Initially(F)$
 $Holds(F, Result(A, S)) :- Holds(F, S), \neg Noninertial(F, A, S)$
 $Holds(alive, Result(shoot, S)) :-$
 $Noninertial(alive, shoot, S) :-$

This program is consistent. Below, $Result[shoot; \dots; shoot, s_0]$ is denoted by $shoot^n$. A Herbrand model of πD_2 is given by the set:

$$\{Holds(alive, shoot^n), Noninertial(alive, shoot, shoot^n) \mid n > 0\}$$

In this model, the e-proposition *shoot* **causes** *alive* overrules the contradicting rule *shoot* **causes** $\neg$ *alive*. π is not complete since D_2 entails all v-propositions, while πD_2 does not.

When D is inconsistent but e-consistent, then πD is inconsistent too. When D is not e-consistent, then π is incomplete iff πD is consistent. Even in such a case, it is often possible to restore the equivalence between D and πD by extending π as follows. For each e-proposition a **causes** f **if** $p_1, \dots, p_n$ with f a positive literal, we add the rule:

$$Initiates(a, f, S) :- Holds(p_1, S), \dots, Holds(p_n, S)$$

For each e-proposition a **causes** $\neg f$ **if** $p_1, \dots, p_n$ with f a positive literal, we add the rule:

$$Terminates(a, f, S) :- Holds(p_1, S), \dots, Holds(p_n, S)$$

In addition, we add the FOL axiom:

$$\forall A, F, S : \neg Initiates(A, F, S) \vee \neg Terminates(A, F, S)$$

Example In πD_2 we have two additional rules:

$$\begin{aligned} Initiates(shoot, alive, S) &:- \\ Terminates(shoot, alive, S) &:- \end{aligned}$$

It is trivial that the resulting program is inconsistent with the FOL axiom.

In some interesting situations, this solution does not work. Consider the following example:

Example The domain D_3 is about flipping a (light) switch. There is one action: *switch* and two fluents, *on* and *off*.

switch causes on if off
switch causes off if on
switch causes $\neg on$ if on
switch causes $\neg off$ if off
Initially on
Initially $\neg off$

D_3 is not consistent, because $\Phi(\text{switch}, \{\text{on}, \text{off}\})$ is not defined consistently. However, starting from the initial situation in which *on* is true and *off* is false and applying *switch* consecutively flips the state of *on* and *off* in such a way that *on* and *off* are never true in the same state. Hence, from the initial state, the problematic state $\{\text{on}, \text{off}\}$ can never be reached. For this reason, πD_3 is consistent, even with *Terminating/3* and *Initiating/3*. Below a Herbrand model is given:

{ *Initially(on)*
Holds(on, switchⁿ) for each even n
Holds(off, switchⁿ) for each odd n
Noninertial(on, switch, switchⁿ) for each n
Noninertial(off, switch, switchⁿ) for each n

$$\begin{array}{ll}
\textit{Initiates}(\textit{switch}, \textit{on}, \textit{switch}^n) & \text{for each odd } n \\
\textit{Terminates}(\textit{switch}, \textit{on}, \textit{switch}^n) & \text{for each even } n \\
\textit{Initiates}(\textit{switch}, \textit{off}, \textit{switch}^n) & \text{for each even } n \\
\textit{Terminates}(\textit{switch}, \textit{off}, \textit{switch}^n) & \text{for each odd } n \quad \}
\end{array}$$

For this example, the semantics of D_3 and πD_3 differ. Which semantics is to be preferred? This is a matter of taste, but intuitively we find the domain description D_3 a sensible theory, and the model a sensible model of the theory. By considering D_3 inconsistent, the semantics of $\mathcal{A}$ is to our taste too severe².

A final example illustrates why v-propositions are added as FOL axioms and not as program clauses.

Example Take the domain D_4 :

$$\begin{array}{l}
a \text{ causes } \neg f \\
f \text{ after } a
\end{array}$$

Obviously this domain is inconsistent. πD_4 is also inconsistent: indeed the completed definition of $\textit{Holds}/2$ subsumes:

$$\textit{Holds}(f, \textit{Result}(a, s_0)) \leftrightarrow \textit{false}$$

That contradicts with the FOL axiom $\textit{Holds}(f, \textit{Result}(a, s_0))$.

On the other hand, adding $\textit{Holds}(f, \textit{Result}(a, s_0))$ as a program clause has the effect of adding the disjunct $F = f \wedge T = \textit{Result}(a, s_0)$ to the completed definition of $\textit{Holds}$. The resulting theory is consistent and has the model:

$$\{\textit{Holds}(f, \textit{Result}(a, s_0)), \textit{Noninertial}(f, a, t_s) | t_s \in \textit{HU}(\mathcal{L}_{D_4})\}$$

4 Reasoning on open logic programs

Traditionally, open programs have been associated with abduction as procedural paradigm. We show how other important procedural paradigms such as deduction and proving satisfiability are feasible and can be emulated by a suitable abduction procedure.

In [6] we proposed SLDNFA, an abductive procedure for normal abductive programs. The soundness of SLDNFA has been proven. Two completeness results were proven: if an SLDNFA execution terminates with failure on a query $\leftarrow Q$, then $\textit{comp}_a(P_D) \models \forall(\neg Q)$; hence no abductive solutions exist. If an SLDNFA execution terminates and generates abductive solutions $\Delta_1, \dots, \Delta_m$, then for each other abductive solution Δ , there exists a Δ_i which is *more general* than Δ in the sense that a skolem substitution θ exists such that $\theta(\Delta_i) \subseteq \Delta$. If Δ_i contains no skolem constants, this simply means that $\Delta_i \subseteq \Delta$ ³.

²Notice that the inconsistency of D_3 can easily be repaired by dropping the fluent *off* and replacing it everywhere by $\neg \textit{on}$. It is unclear to us whether such a solution exists in general when the semantics of D and πD differ.

³Any other abductive procedure which satisfies the same completeness results can be used. At the time of writing, SLDNFA seems to have advantages over other published procedures in two senses: it suffers less from floundering since it does not flounder on abducible atoms and more powerful completeness results have been proved for it.

SLDNFA is not developed for dealing with FOL axioms, but there is a general technique to transform an open logic program P and a theory of FOL axioms $\mathcal{T}$ into an equivalent open logic program P' . This transformation technique is a trivial extension of the transformation proposed in [23]. In a first step, P is extended with:

$$false \text{ :- } \neg F$$

for each FOL axiom $F \in \mathcal{T}$. The result is an open logic program with *general clauses*. In the second step, it is transformed to a normal open logic program P' using the technique in [23]. The transformation is correct in the sense that $P + \mathcal{T}$ is equivalent with $P' + \{\neg false\}$ according to completion semantics. The remaining FOL axiom $\neg false$ can be added as an extra literal to the query to be solved by the abductive solver. This result shows that FOL (with *FEQ*) and open logic programming have the same expressivity in the strongest possible sense, namely on the level of logical equivalence. The proof of this result is an extension of the proof in [23] and can be found in [4].

For a domain description D , the transformation of the FOL axioms of πD to an open logic program is trivial. A ground atom $Holds(f, Result[a_1; \dots; a_n, s_0])$ is transformed to:

$$false \text{ :- } \neg Holds(f, Result[a_1; \dots; a_n, s_0])$$

A ground negative literal $\neg Holds(f, Result[a_1; \dots; a_n, s_0])$ is transformed to:

$$false \text{ :- } Holds(f, [a_1; \dots; a_n, s_0])$$

Applying this technique on the FOL axioms of the Murder Mystery domain D_{MM} , we obtain an open program P' , in which the following rules:

$$\begin{aligned} false &\text{ :- } \neg Holds(alive, s_0) \\ false &\text{ :- } Holds(alive, [shoot; wait, s_0]) \end{aligned}$$

are substituted for the FOL axioms of πD_{MM} .

An abductive procedure generates explanations for a given observation on the problem domain. Here we can take $\neg false$ as an observation. SLDNFA solves the query $\leftarrow \neg false$ and returns the solution:

$$\Delta_1 = \{Initially(loaded), Initially(alive)\}$$

This not only gives an explanation for $\neg false$, but also proves that πD_{MM} is consistent. The precise reasoning goes as follows: the resulting program $P' + \Delta$ is an acyclic program and therefore is consistent wrt completion semantics [1]. $P' + \Delta$ comprises $P_{D_{MM}}$ and entails both $\neg false$ and $\neg false \leftrightarrow IC_{D_{MM}}$. Hence a model of $P' + \Delta$ is a model of πD .

An abductive procedure can also be used for deduction. For example, we want to prove that $\pi D_{MM} \models Initially(loaded)$ or equivalently that the theory $\pi D_{MM} + \neg Initially(loaded)$ is inconsistent. To prove that, we add the extra rule :

$$false \text{ :- } Initially(loaded)$$

Now SLDNFA fails finitely on the query $\leftarrow \neg false$. From the first completeness result of SLDNFA, it follows that $\pi D_{MM} + \neg Initially(loader)$ is inconsistent. Notice that a completeness result for abduction is used here as a soundness result for deduction.

An abductive procedure allows reasoning under uncertainty. By dropping the v-proposition **Initially alive** from the Murder Mystery domain D_{MM} , an incomplete domain description D'_{MM} is obtained. Using $\pi D'_{MM}$, SLDNFA answers the goal $\leftarrow \neg false$ by returning the answer $\Delta_2 = \{\}$. The original solution Δ_1 is still a solution but is not generated. This does not conflict with the completeness result of SLDNFA because $\Delta_2 \subset \Delta_1$.

That we have uncertainty in this domain description becomes obvious when we want to know whether **Initially alive** is *possible* according to D'_{MM} . This is done by posing the query $\leftarrow \neg false, Initially(alive)$. SLDNFA proves that **Initially alive** is possible by returning Δ_1 .

Deduction under uncertainty is possible. Observe that D'_{MM} entails:

$$\mathbf{Initially} \neg alive \vee \mathbf{Initially} loaded$$

SLDNFA can prove this. This is done by transforming the negation of the disjunction to:

$$\begin{aligned} false &:- Initially(loader) \\ false &:- \neg Initially(alive) \end{aligned}$$

After adding these rules to $\pi D'_{MM}$, SLDNFA fails finitely on $\leftarrow \neg false$. This proves the disjunction.

The above experiments show in the first place that though open/ abductive logic programs are traditionally associated with abduction as procedural paradigm, other procedural paradigms such as deduction and consistency proving are of interest. This illustrates our argument that an *abductive program* is better called an *open program*. In the second place, the experiments show that a suitable abductive procedure can be used to emulate these paradigms.

5 The Gelfond & Lifschitz approach

We recall the transformation proposed in [17], from $\mathcal{A}$ domain descriptions to extended logic programs. For any domain description D , $\pi_{GL}D$ is defined as the extended logic program containing the following extended clauses:

- Four inertia rules:

$$(5.1) \quad Holds(F, Result(A, S)) :- Holds(F, S), not Noninertial(F, A, S)$$

$$(5.2) \quad \neg Holds(F, Result(A, S)) :- \neg Holds(F, S), not Noninertial(F, A, S)$$

$$(5.3) \quad Holds(F, S) :- Holds(F, Result(A, S)), not Noninertial(F, A, S)$$

$$(5.4) \quad \neg Holds(F, S) :- \neg Holds(F, Result(A, S)), not Noninertial(F, A, S)$$

- Each v-proposition f **after** $a_1; \dots; a_n$, is transformed into:

$$(5.5) \quad \text{Holds}(f, \text{Result}[a_1; \dots; a_n, s_0]) :-$$

Recall that $\text{Holds}(\neg f, \dots)$ denotes $\neg \text{Holds}(f, \dots)$.

- Each e-proposition a **causes** f **if** $p_1, \dots, p_n$ is translated into $2n+2$ rules. Below, $\overline{\text{Holds}(f, S)}$ denotes the complement of $\text{Holds}(f, S)$ with respect to $\neg$.

$$(5.6) \quad \text{Holds}(f, \text{Result}(a, S)) :- \text{Holds}(p_1, S), \dots, \text{Holds}(p_n, S)$$

$$(5.7) \quad \text{Noninertial}(|f|, a, S) :- \text{not } \overline{\text{Holds}(p_1, S)}, \dots, \text{not } \overline{\text{Holds}(p_n, S)}$$

For each i , $1 \leq i \leq n$:

$$(5.8) \quad \text{Holds}(p_i, S) :- \overline{\text{Holds}(f, S)}, \text{Holds}(f, \text{Result}(a, S))$$

$$(5.9) \quad \overline{\text{Holds}(p_i, S)} :- \overline{\text{Holds}(f, \text{Result}(a, S))}, \\ \text{Holds}(p_1, S), \dots, \text{Holds}(p_{i-1}, S), \\ \text{Holds}(p_{i+1}, S), \dots, \text{Holds}(p_n, S)$$

[17] gives the intuition behind the translation and gives a soundness theorem for all domain descriptions D provided D does not contain *similar* e-propositions, i.e. e-propositions which only differ by the preconditions. A comparison of $\pi_{GL}D$ with πD is of interest. Observe that if the two negations *not* and $\neg$ in extended logic programs are mapped both on " $\neg$ " in open programs, we find the following correspondences between $\pi_{GG}D$ and πD : (5.1) $\leftrightarrow$ (3.2), (5.5) $\leftrightarrow$ (3.5), (5.6) $\leftrightarrow$ (3.3) (if f is a positive literal), (5.7) $\leftrightarrow$ (3.4), while (5.2), (5.3), (5.4), (5.8) and (5.9) lack in πD .

A striking fact is that $\pi_{GL}D$ contains four inertia rules instead of one in πD . (5.1) and (5.2) are *forward persistence rules* for respectively positive and negative fluents. (5.3) and (5.4) are *backward persistence rules* for again positive and negative fluents. Clearly (5.2), (5.3) and (5.4) are natural rules, which are expected to hold in any correct formalisation. Therefore, they must be subsumed by πD , otherwise π could never be sound and complete. As a matter of fact, it is straightforward to prove that for each of the extended rules in $\pi_{GL}D$, the corresponding clause is subsumed by $\text{comp}_a(P_D)$, where P_D is the logic program part of πD . For example, notice that from the classical logic point of view the rules (5.1) and (5.4) are equivalent and so are the rules (5.2) and (5.3). This immediately gives that $\text{comp}_a(P_D)$ subsumes (5.4). Clauses corresponding to (5.2) and (5.3) can be derived from the completed definition (6.6) of $\text{Holds}/2$ in P_D . Substitute $\text{Result}(A, S)$ for T . After simplification one obtains:

$$\forall F, A, S : \text{Holds}(F, \text{Result}(A, S)) \leftrightarrow E_1 \vee \dots \vee E_n$$

where E_1 is of the form:

$$\text{Holds}(F, S) \wedge \neg \text{Noninertial}(F, A, S)$$

and for each e-proposition a **causes** f **if** $p_1, \dots, p_n$ with f is a positive literal, there is an E_i of the form:

$$F = f \wedge A = a \wedge \text{Holds}(p_1, S) \wedge \dots \wedge \text{Holds}(p_n, S)$$

Now, it is easy to see that $\text{comp}_a(P_D)$ satisfies:

$$\begin{aligned} \forall F, A, S : F = f \wedge A = a \wedge \text{Holds}(p_1, S) \wedge \dots \wedge \text{Holds}(p_n, S) \\ \rightarrow \text{Noninertial}(F, A, S) \end{aligned}$$

By dropping $\neg \text{Noninertial}(F, A, S)$ from the first disjunct and substituting $\text{Noninertial}(F, A, S)$ for the other disjuncts, we find:

$$\forall F, A, S : \text{Holds}(F, \text{Result}(A, S)) \rightarrow \text{Holds}(F, S) \vee \text{Noninertial}(F, A, S)$$

Simple rewriting gives formulas corresponding to (5.2) and (5.3).

A shortcoming of π_{GL} is its incompleteness. [17] gives the following example D_5 :

a **causes** f **if** f
 f **after** a

Clearly D_5 entails **Initially** f . However, $\text{Initially}(f)$ is not entailed by $\pi_{GL}D_5$. On the other hand, notice that D_5 is e-consistent. Therefore, πD_5 implies $\text{Initially}(f)$ and SLDNFA can prove that.

Another problem of π_{GL} shows up when $\mathcal{A}$ is extended to allow predicates. Consider the following rule:

$$\text{Pick}(X, \text{Obj}) \text{ **causes** } \text{thief}(X) \text{ **if** } \text{owner}(Y, \text{Obj}), X \neq Y$$

which says that X becomes a thief if he picks an object Obj which he does not own. The translation to open logic programs does not require any modification. π produces:

$$\begin{aligned} \text{Holds}(\text{thief}(X), \text{Result}(\text{Pick}(X, \text{Obj}), S)) &:- \text{Holds}(\text{owner}(Y, \text{Obj}), S), \\ &X \neq Y \end{aligned}$$

$$\text{Noninertial}(\text{thief}(X), \text{Pick}(X, \text{Obj}), S) :- \text{Holds}(\text{owner}(Y, \text{Obj}), S), X \neq Y$$

For π_{GL} , there are problems with the rules of type (5.8):

$$\begin{aligned} \text{Holds}(\text{owner}(Y, \text{Obj}), S) &:- \neg \text{Holds}(\text{thief}(X), S), \\ &\text{Holds}(\text{thief}(X), \text{Result}(\text{Pick}(X, \text{Obj}), S)) \end{aligned}$$

$$X \neq Y :- \neg \text{Holds}(\text{thief}(X), S), \text{Holds}(\text{thief}(X), \text{Result}(\text{Pick}(X, \text{Obj}), S))$$

These rules say that when X becomes thief by picking something in situation S , then each Y is owner at situation S and no Y is equal to X . This is a contradiction. The problem is that Y should not be universally but existentially quantified. The following formulas are subsumed by πD but are not extended clauses:

$$\begin{aligned} \forall \text{Obj}, S, X : \exists Y : \text{Holds}(\text{owner}(Y, \text{Obj}), S) &\leftarrow \neg \text{Holds}(\text{thief}(X), S), \\ &\text{Holds}(\text{thief}(X), \text{Result}(\text{Pick}(X, \text{Obj}), S)) \\ \forall \text{Obj}, S, X : \exists Y : X \neq Y &\leftarrow \neg \text{Holds}(\text{thief}(X), S), \\ &\text{Holds}(\text{thief}(X), \text{Result}(\text{Pick}(X, \text{Obj}), S)) \end{aligned}$$

The translation π to open logic programs is superior to the translation π_{GL} to extended logic programs. π_{GL} creates a higher number of rules, is incomplete, suffers from problems with similar e-propositions and is not directly extendible to the predicate case. The open program approach seems more understandable because only one negation occurs, is sound even with similar e-propositions, is complete for all reasonable domain descriptions and applies without modification for the predicate case. Proofs are easy compared with the proofs in [17].

6 Dung's approach

Independent from the work presented here, another approach, quite similar to ours in a number of aspects, has been developed for translating $\mathcal{A}$ domain descriptions to a logic program formalism. This work, described in [11], maps $\mathcal{A}$ to the logic program formalism and defines a special purpose semantics, similar to the completion semantics.

On the syntactical level, the most important difference with our approach is the symmetrical treatment of fluent symbols f and their negation $\neg f$. The translation $\pi_{Du}D$ contains our frame axiom (6.2), and contains for each e-proposition a **causes** f if $p_1, \dots, p_n$ (f positive or negative) the following rules:

$$Holds(f, Result(a, S)) :- Holds(p_1, S), \dots, Holds(p_n, S)$$

$$Noninertial(f, a, S) :- Holds(p_1, S), \dots, Holds(p_n, S)$$

$$Noninertial(f^*, a, S) :- Holds(p_1, S), \dots, Holds(p_n, S)$$

Here f^* denotes the complement of f . Contrary to our approach, here fluent literals like $\neg f$ appear within $Holds/2$. Each v-proposition f **after** $a_1; \dots; a_n$ is transformed to the denial:

$$\leftarrow Holds(f^*, Result(a_1; \dots; a_n, s_0))$$

In addition, for each fluent symbol f the following constraints are added:

$$\leftarrow Holds(f, S), Holds(\neg f, S)$$

$$Holds(\neg f, s_0) \leftrightarrow \neg Holds(f, s_0)$$

These additional constraints are necessary due to the symmetrical treatment of a fluent f and its negation $\neg f$. In contrast, our transformation maps $\mathcal{A}$ domains to a more classical situation calculus, which treats the positive and negative literals asymmetrically. Considering only Dung's solution, one might wonder whether the symmetrical treatment is crucial for solving the frame problem in logic programming or whether it has advantages. Our transformation shows that a symmetrical treatment is not necessary and leads to substantially more rules.

The semantics of $\pi_{Du}D$ is defined via a domain dependent variant of the completion semantics. It contains FEQ and the normal completed definition of $Noninertial/3$, but a specialised version of the completed definition of $Holds/2$. $\pi_{Du}D$ does not contain rules with $Holds(F, s_0)$ in the head. As a consequence the standard completion would imply that

$$\forall F : \neg \text{Holds}(F, s_0)$$

[11] avoids this by the following alternative:

$$\forall F, A, T : \text{Holds}(F, \text{Result}(A, S)) \leftrightarrow E_1 \vee \dots \vee E_n$$

$$\begin{aligned} \text{with: } E_1 &= \text{Holds}(F, S) \wedge \neg \text{Noninertial}(F, A, S) \\ E_i (i > 1) &= F = f \wedge A = a \wedge \text{Holds}(p_1, S) \wedge \dots \wedge \text{Holds}(p_n, S) \end{aligned}$$

such that there is a correspondence between the $E_i (i > 1)$ and the e-propositions ***a causes f if $p_1, \dots, p_n$*** (f positive or negative). This formula says nothing about $\text{Holds}(F, s_0)$, and therefore, we get a similar semantics as in our approach, but without *Initially/1*.

Dung extends $\mathcal{A}$ to the predicate case and gives an application for integrity checking of a database update. For this purposes, a predicate domain description D is developed to represent the effects of primitive and compound update operations. An unfolding partial evaluation procedure with constructive negation can then be applied on $\pi_{Du}(D)$ to check the consistency of one or more integrity constraints in the database after the update. The use of a similar unfolding procedure has been proposed earlier in the context of abductive logic programming by [3]. As Dung here, Console et al. use a general unfolding procedure to reduce a given query by unfolding all defined literals until a first order logic formula is obtained, containing only abductive literals. The resulting formula is equivalent with the original query and the equivalence with in the left the query and in the right the computed formula is called the explanation formula. Though such a procedure can be used to solve abductive problems, it is much weaker than an abductive procedure like SLDNFA because naive unfolding cannot guarantee the consistency of the right-hand of the explanation formula. In [5, 9], we show that SLDNFA can be interpreted also as a kind of an unfolding procedure, but SLDNFA is equipped with a mechanism for checking the consistency of each of the disjuncts at the right hands of the generated explanation formula. In our opinion, this is absolutely necessary because even for simple problems, an explanation formula can be so complex that it is impossible for humans to grasp its meaning and to check its consistency.

An interesting statement in [11] is related to the following example D_6 (a syntactical simplification of the example in theorem (8) in [11]):

$$\begin{aligned} &\text{Initially } f \\ &a \text{ causes } \neg f \text{ if } g \end{aligned}$$

Dung observes that the Gelfond and Lifschitz transformation $\pi_{GL}D_6$ has two answer sets:

$$\begin{aligned} Z_1 &= \{\text{Holds}(f, s_0)\} \cup \{\text{Noninertial}(f, a, a^n) \mid n \in \mathbb{N}\} \\ Z_2 &= \{\text{Holds}(f, a^n), \neg \text{Holds}(g, a^n) \mid n \in \mathbb{N}\} \end{aligned}$$

In Z_1 , $\text{Holds}(f, a^n)$ ($n > 0$) and $\text{Holds}(g, a^n)$ ($n \geq 0$) are unknown since neither the atom nor its negation appears in Z_1 . Z_2 corresponds to a two-valued model, obtained by having f initially true and g false, a situation which is preserved when applying a . Then Dung argues that "*it is obvious that only the first solution captures the intended semantics of D_6 , for if we don't know anything about g , it is impossible*

to say anything about the outcome of a ". Remarkable now is that Z_2 corresponds to a model of Dung's $\pi_{D_u}D_6$.

Here we touch a fundamental issue in knowledge representation: how should incomplete knowledge be formalised in a logic? What Dung suggests in this statement is that a model of a theory should reflect what one knows about the world: a model should be the set of atoms which are known to be true. To formalise this view, one needs 3-valued interpretations, in which atoms are either true, unknown or false. Though technically spoken an answer set is 2-valued, it encodes a 3-valued interpretation due to the presence of both a predicate and its explicit negation: given an answer set M , $p \in M$ means that p is true; $\neg p \in M$ means that p is false; $p \notin M$ and $\neg p \notin M$ means that p is unknown. The fourth case that both p and $\neg p$ occur in M is excluded by definition of answer set. Under the view of an interpretation as a 3-valued representation of beliefs, Z_1 is indeed the intended 3-valued model and Z_2 is incorrect, since we cannot know for sure that g is not true initially and hence f is possibly terminated by applying a .

The above view on interpretations is not the classical logic view on incomplete knowledge. In the classical logic view, a model is a mathematical abstraction of a possible state of the world, not the set of ground atoms which are known by or provable from the theory. To have incomplete knowledge on g in the initial state, is reflected by the fact that there are models in which g is initially true and others in which g is initially false.

Of the two views, the classical view is definitely the richest one. Indeed, consider the following formula:

$$\neg \text{Holds}(g, s_0) \rightarrow \text{Holds}(f, \text{Result}(a, s_0))$$

which expresses that if g does not hold initially, then f holds after applying a on the initial state. Surely one will agree that this conclusion is correct. It can be proven from π_{D_6} and $\pi_{D_u}D_6$ using a simple case analysis: due to the 2-valued nature of classical models, in any model either g is true at s_0 , or g is false at s_0 . In the first case, the condition of the implication is not satisfied. In the second case, the inertia rule in π_{D_6} and $\pi_{D_u}D_6$ can be applied to obtain that f holds at $\text{Result}(a, s_0)$.

On the other hand, the formula is not true in the answer set Z_1 , and hence is not implied by $\pi_{GL}D_6$. Due to the 3-valued nature of Z_1 , the case analysis which holds for π and π_{D_u} , does not apply in Z_1 , since there the third case holds: neither $\text{Holds}(g, s_0)$ nor $\neg \text{Holds}(g, s_0)$ is true. Due to this, $\text{Noninertial}(f, a, s_0)$ can be derived, and hence, the inertia rule is disabled. Hence, $\text{Holds}(f, \text{Result}(a, s_0))$ cannot be derived.

The 3-valued nature of answer set semantics is the cause of the incompleteness of π_{GL} . For example, take the example $\pi_{GL}(D_5)$ which fails to entail the desired conclusion $\text{Holds}(f, s_0)$. It is straightforward to see that in any answer set in which $\text{Holds}(f, \text{Result}(a, s_0))$ is true, then $\neg \text{Holds}(f, s_0)$ is inconsistent with rule (5.8). In a 2-valued semantics, this would immediately entail that $\text{Holds}(f, s_0)$ is true. However, $\pi_{GL}(D_5)$ has a unique answer set in which $\text{Holds}(f, s_0)$ is unknown, i.e. neither $\text{Holds}(f, s_0)$ nor $\neg \text{Holds}(f, s_0)$ holds.

Actually, it has been argued before by Moore [27] that the form of reasoning by 2 cases (i.e. something is either true or false) is crucial for reasoning on uncertainty. This reasoning principle is invalidated in 3-valued semantics. It is often believed that logic programs under a 3-valued semantics (like well-founded semantics) can be used

to represent uncertainty. Generalising the above observations, we believe that only very restricted forms of uncertainty can be represented in any 3-valued logic.

Though Dung does not investigate this, his technique for representing uncertainty can be generalised to other applications in a rather straightforward way. The resulting formalism would be a generalisation of the abductive logic program formalism, in which it is possible to give definitions at the level of atoms instead of predicates. Such a generalised abductive program might be seen as a tuple of a set of defined atoms $D = \{p_1(\bar{t}_1), \dots, p_n(\bar{t}_n)\}$ and a set of normal clauses such that if two atoms $p(\bar{t}), p(\bar{s})$ occur in D , then they do not unify and if $p(\bar{s}) :- L_1, \dots, L_n$ belongs to the program, then $p(\bar{s})$ should be an instance of an atom $p(\bar{t}) \in D$. Such a clause can be viewed as a defining clause for $p(\bar{t})$.

The semantics of such a generalised open program could then be given by *FEQ* and for each defined atom $p_i(\bar{t}_i)$ a completed definition $p_i(\bar{t}_i) \leftrightarrow \dots$ constructed using its defining clauses.

However, this generalisation does not produce extra expressivity, as the formalism can be easily and elegantly translated to the open logic program formalism. Given such a generalised logic program P , we can translate P to an open logic program P' obtained by adding for each predicate p/n one new undefined predicate *undefined* $_p/n$ and taking the clauses of P and adding for each predicate p/n with defined atoms $p(\bar{t}_1), \dots, p(\bar{t}_n)$ one new clause:

$$p(\bar{X}) :- \neg(\exists \bar{Y}_1. \bar{X} = \bar{t}_1), \dots, \neg(\exists \bar{Y}_n. \bar{X} = \bar{t}_n), \text{undefined}_p(\bar{X})$$

Here $\bar{Y}_i$ is the set of variables $\text{var}(\bar{t}_i)$ and $\bar{X} = \bar{t}_i$ denotes the conjunction of equality atoms $X_j = t_j^i$. The completion of P' can easily be proven to be a conservative extension of P .

Though our approach was developed independently of Dung's approach, it clearly shows similarities it. Both generate a form of situation calculus using the logic program syntax, and the semantics of the resulting programs is given by first order logic theories, constructed using different forms of completion of sets of clauses.

Despite the similarities, we believe that our work contributes in several important aspects. Recall our two main goals of this experiment with $\mathcal{A}$: first, to show the role of open/abductive logic programs as a general purpose logic for representing uncertainty and second, to show the role of a suitable abductive procedure (as SLDNFA) for solving different computational problems, including deduction under uncertainty, abduction or the generation of explaining hypotheses and consistency proving. To both goals, Dung's paper does not really contribute.

With respect to the first goal, Dung proposes a novel formalism. He does not relate this formalism to open/abductive logic programming. As argued in section 3, we could prove the correctness of the transformation π not only wrt completion semantics but wrt 3-valued completion semantics, generalised stable model semantics, generalised well-founded semantics, justification semantics, etc. This aspect is necessarily lacking in Dung's paper, since he doesn't consider abductive logic programming in the first place.

With respect to the second goal, we have shown how SLDNFA can be used to solve three important types of computational problems:

- deduction (under uncertainty)

- abduction
- consistency proving

Dung does not consider these classical forms of reasoning but focusses on the use of an unfolding partial evaluation procedure for integrity checking of database updates. The question whether such an unfolding procedure is of practical use for general problem solving in temporal domains is not dealt with in Dung's paper.

7 Discussion

In [12], a First Order Logic solution to the frame-problem was proposed. [30] uses the same type of theory to formalise database evolution. This type of theory is a form of situation calculus which shows a strong similarity with the completion of a program πD . *Result*/2 is replaced by *do*/2. Instead of using the meta predicate *Holds*/2, [30] adds one additional argument to each fluent predicate; i.e. an atom *Holds*($p(x), t$) is contracted to the atom $p(x, t)$. As a consequence the law of inertia has to be stated for each fluent. An example taken from [30] is given below:

$$\begin{aligned} \forall St, C, A, S : Poss(A, S) \rightarrow (enrolled(St, C, do(A, S)) \leftrightarrow \\ A = register(St, C) \vee \\ enrolled(St, C, S) \wedge A \neq drop(A, C)) \end{aligned}$$

The rule says that when action A may be executed in situation S ($Poss(A, S)$), then student St is enrolled in course C at time $do(A, S)$ iff A is an action of registering St in C or, St was enrolled at S and A is not an action of dropping St from the course C .

As said above, this type of theory shows a strong similarity with the completion of the open programs produced by π and with the completions of the logic programs found in [14] and [1]. If we forget about $Poss(A, S)$, and introduce *Result* and *Holds*, we find:

$$\begin{aligned} \forall St, C, A, S : Holds(enrolled(St, C), Result(A, S)) \leftrightarrow \\ A = register(St, C) \vee \\ Holds(enrolled(St, C), S) \wedge A \neq drop(St, C) \end{aligned}$$

Similar formulas are subsumed by πD . The first disjunct corresponds to a rule initiating *enrolled*(St, C) by *enregister*(St, C). The second disjunct corresponds to the law of inertia, with $\neg Noninertial(enrolled(St, C), A, S)$ replaced by its definition: *drop*(St, C) is the only action which terminates *enrolled*(St, C)⁴.

It is remarkable and frustrating that the above monotonic solution in classical logic has not been discovered much earlier. Some people have been experimenting with situation calculus in logic programming ever since the first experiments in Kowalski's *Logic for problem solving* [21] (first edition in 76). That Prolog with negation as finite failure can be interpreted as a sound theorem prover wrt to the completion of a Prolog program is known since Clark's work [2] in 78. When [18] introduced the YTS problem to show the failure of several nonmonotonic solutions to the frame problem,

⁴Here the version of *Noninertial*/3 is needed which contains only rules for terminating effects of actions.

it was soon realised by some in the logic programming community that the Prolog solution behaved perfectly correct. However, if Prolog can prove that *waiting* does not unload the gun and that the turkey is *dead* after *loading*, *waiting* and *shooting*, then from Clark's soundness theorem it follows immediately that the completion of the Prolog program entails these conclusions also, and hence that the completion of the Prolog program provides a (monotonic) solution to the frame problem. Generalising this example, it is a relatively simple step to find that the principle of completion of implications gives an elegant and general solution to the frame problem. Yet, it seems that for many years, nobody, neither in the logic programming community nor in the A.I. community has come to this obvious conclusion or was interested in it.

In the past, another approach has been explored for temporal reasoning, based on event calculus [22]. [13] and [31] have simplified event calculus and have extended it with abduction for the purpose of planning. [31] extended event calculus to deal with necessary preconditions of actions. [25] implemented a planning system based on this formalism. Other work has been done to extend event calculus with continuous actions [32] and time granularity [15], [26]. Recently [10] applied abductive event calculus to solve a number of benchmark problems in temporal reasoning, such as the Murder Mystery, the Stolen Car problem, the Walking Turkey Shooting problem and the Russian Turkey Shooting problem. The latter problem contains an indeterminate action. Situation and event calculus seem two non-equivalent ways of representing time and action. $\mathcal{A}$ domain descriptions cannot (easily) be translated to event calculus, because $\mathcal{A}$ assumes a *branching time* philosophy, as in the situation calculus, whereas event calculus assumes a linear time philosophy. A deep analysis of situation versus event calculus is beyond the scope of the paper.

The $\mathcal{A}$ language is designed to represent one form of incomplete temporal knowledge: on the initial situation. A totally different form of incomplete knowledge appears with indeterminate actions. Our transformation from $\mathcal{A}$ to open programs can easily be adapted in order to deal with such actions, which shows again the expressivity of the open program formalism for representing uncertainty. It turns out that the technique used in [10] to represent indeterminate actions in the context of event calculus can easily be translated to situation calculus. The Russian Turkey Shooting problem is a variant of the Yale Turkey Shooting problem in which one additional action *spinning* of spinning the gun's chamber occurs. The effect is that the gun is possibly unloaded. Below we allow e-propositions of the form *a possibly causes f if $p_1, \dots, p_n$* . The problem is formalised as follows:

Initially alive
Initially loaded
load causes loaded
shoot causes $\neg$ alive if loaded
shoot causes $\neg$ loaded
spinning possibly causes $\neg$ loaded

The semantics of $\mathcal{A}$ should be adapted. While in $\mathcal{A}$, a successor state is completely determined by the action and the previous state, this is not the case with indeterminate actions. Therefore, in the extended version the transition function should be replaced by a transition relation. While it is beyond the scope of this paper to work out this semantics in detail, it is easy to show how the transformation could

be adapted to model this kind of indeterminate actions. The indeterminism can be captured by introducing an undefined *GoodLuck/2* predicate:

$$\text{Noninertial}(\text{loaded}, \text{spinning}, S) \text{ :- } \text{GoodLuck}(\text{spinning}, S)$$

The above clause has the effect that the rule of inertia is disabled for *loaded* iff good luck occurs at the spinning action in state *S*. In general, for each clause *a possibly causes f* if $p_1, \dots, p_n$ the following rule must be introduced:

$$\text{Noninertial}(|f|, a, S) \text{ :- } \text{Holds}(p_1, S), \dots, \text{Holds}(p_n, S), \text{GoodLuck}(a, S)$$

For a positive *f*, in addition the following rule is added:

$$\text{Holds}(f, \text{Result}(a, S)) \text{ :- } \text{Holds}(p_1, S), \dots, \text{Holds}(p_m, S), \text{GoodLuck}(a, S)$$

8 Summary

We presented a sound and complete transformation π from $\mathcal{A}$ domains to open logic programs with FOL axioms. We have illustrated the use of SLDNFA for abductive and deductive reasoning and satisfiability proving under uncertainty. The transformation of Gelfond and Lifschitz is more complex, is not complete due to the 3-valued nature of answer set semantics, is only sound for domains without e-similar actions and cannot be extended to the predicate case (at least not without imposing other syntactic constraints). Moreover, no reasoning procedure is currently described for the resulting programs. Dung's approach is in some aspects similar to ours and provides a reasoning procedure, but is still more complex than ours, has the disadvantage of relying on a special purpose logic and does not show the application of the reasoning procedure for classical forms of reasoning such as deduction, abduction and consistency proving.

We have investigated also a number of typical temporal reasoning issues. Although in πD only forward persistence axioms are contained, the completion of πD subsumes backward persistence axioms. We have also shown how to extend $\mathcal{A}$ with indeterminate actions.

From a more general perspective, this work can be viewed as a -successful- experiment in the declarative representation of and diverse forms of automated reasoning on incomplete knowledge using open logic programming and an abductive procedure.

Acknowledgements

This research is partially supported by the K.U.Leuven research coordination, by ESPRIT BRA Compulog-II under contract 6810 and by the Belgian National Fund for Scientific Research.

References

- [1] K.R. Apt and M. Bezem. Acyclic programs. In *Proc. of the International Conference on Logic Programming*, pages 579–597. MIT press, 1990.

- [2] K.L. Clark. Negation as failure. In H. Gallaire and J. Minker, editors, *Logic and Databases*, pages 293–322. Plenum Press, 1978.
- [3] L. Console, D. Theseider Dupre, and P. Torasso. On the relationship between abduction and deduction. *Journal of Logic and Computation*, 1(5):661–690, 1991.
- [4] M. Denecker. *Knowledge Representation and Reasoning in Incomplete Logic Programming*. PhD thesis, Department of Computer Science, K.U.Leuven, 1993.
- [5] M. Denecker and D. De Schreye. On the duality of abduction and model generation. In *Proc. of the International Conference on Fifth Generation Computer Systems*, 1992.
- [6] M. Denecker and D. De Schreye. SLDNFA; an abductive procedure for normal abductive programs. In K.R. Apt, editor, *Proc. of the International Joint Conference and Symposium on Logic Programming*, pages 686–700, 1992.
- [7] M. Denecker and D. De Schreye. Justification semantics: a unifying framework for the semantics of logic programs. In *Proc. of the Logic Programming and Nonmonotonic Reasoning Workshop*, pages 365–379, 1993.
- [8] M. Denecker and D. De Schreye. Representing incomplete knowledge in abductive logic programming. In *Proc. of the International Symposium on Logic Programming*, pages 147–163, 1993.
- [9] M. Denecker and D. De Schreye. On the Duality of Abduction and Model Generation in a Framework for Model Generation with Equality. *Journal of Theoretical Computer Science*, 122(1&2):225–262, 1994.
- [10] M. Denecker, L. Missiaen, and M. Bruynooghe. Temporal reasoning with abductive event calculus. In *Proc. of the European Conference on Artificial Intelligence*, 1992.
- [11] P.M. Dung. Representing Actions in Logic Programming and its Applications in Database Updates. In *Proc. of the International Conference on Logic Programming*, 1993.
- [12] C. Elkan. Reasoning about Action in First-Order Logic. In *Proc. of the CSCSI-92*, 1992.
- [13] K. Eshghi. Abductive planning with Event Calculus. In R.A. Kowalski and K.A. Bowen, editors, *Proc. of the International Conference on Logic Programming*, 1988.
- [14] C. Evans. Negation as failure as an approach to the Hanks and McDermott problem. In *Proc. of the second International Symposium on Artificial Intelligence*, 1989.
- [15] C. Evans. The Macro-Event Calculus: Representing Temporal Granularity. In *Proc. of PRICAI, Tokyo*, 1990.

- [16] M. Gelfond and V. Lifschitz. Logic Programs with Classical Negation. In D.H.D. Warren and P. Szeredi, editors, *Proc. of the 7th International Conference on Logic Programming 90*, page 579. MIT press, 1990.
- [17] M. Gelfond and V. Lifschitz. Describing Action and Change by Logic Programs. In *Proc. of the 9th Int. Joint Conf. and Symp. on Logic Programming*, 1992.
- [18] S. Hanks and D. McDermott. Default reasoning, nonmonotonic logic, and the frame problem. In *Proc. of the National Conference on Artificial Intelligence, Philadelphia*, pages 328–333, 1986.
- [19] P. M. Hill and J. W. Lloyd. Analysis of meta-programs. In H. D. Abramson and M. H. Rogers, editors, *Proceedings of Meta88*, pages 23–51. MIT Press, 1989.
- [20] A.C. Kakas and P. Mancarella. Generalised stable models: a semantics for abduction. In *Proc. of the European Conference on Artificial Intelligence*, 1990.
- [21] R.A. Kowalski. *Logic for problem solving*. Elsevier Science Publisher, 1976.
- [22] R.A. Kowalski and M. Sergot. A logic-based calculus of events. *New Generation Computing*, 4(4):319–340, 1986.
- [23] J.W. Lloyd and R.W. Topor. Making prolog more expressive. *Journal of Logic Programming*, 1(3):225–240, 1984.
- [24] B. Martens and D. De Schreye. A perfect Herbrand semantics for untyped vanilla meta-programming. In K.R. Apt, editor, *Proceedings of the Joint International Conference and Symposium on Logic Programming*, pages 511–525, Washington, November 1992. MIT Press.
- [25] L.R. Missiaen. *Localized abductive planning with the event calculus*. PhD thesis, Department of Computer Science, K.U.Leuven, 1991.
- [26] A. Montanari, E. Maïm, E. Ciapessoni, and E. Ratto. Dealing with Time Granularity in the Event Calculus. In *Proc. of the International Conference on Fifth Generation Computer Systems*, pages 702–712, 1992.
- [27] R.C. Moore. The role of logic in knowledge representation and commonsense reasoning. In *Proc. of AAAI-82*, pages 428–433, 1982.
- [28] L.M. Pereira, J.N. Aparicio, and J.J. Alferes. Hypothetical Reasoning with Well Founded Semantics. In B. Mayoh, editor, *Proc. of the 3th Scandinavian Conference on AI*. IOS Press, 1991.
- [29] C.S. Pierce. *Philosophical Writings of Pierce*. Dover Publications, New York, 1955.
- [30] R. Reiter. Formalizing Database Evolution in the Situation Calculus. In *Proc. of the International Conference on Fifth Generation Computer Systems*, pages 600–609, 1992.
- [31] M. Shanahan. Prediction is deduction but explanation is abduction. In *Proc. of the IJCAI89*, page 1055, 1989.

- [32] M. Shanahan. Representing continuous change in the event calculus. In *Proc. of the European Conference on Artificial Intelligence*, page 598, 1990.