

A Qualitative Model for Temporal Reasoning with Incomplete Information

Hector Geffner*

Departamento de Computación
Universidad Simón Bolívar
Aptdo 89000, Caracas 1080-A, Venezuela

Abstract

We develop a qualitative framework for temporal reasoning with incomplete information that features a modeling language based on rules and a semantics based on infinitesimal probabilities. The framework relates logical and probabilistic models, and accommodates in a natural way features that have been found problematic in other models like non-determinism, action qualifications, parallel actions, and abduction to actions and fluents.

Introduction

Logic, probabilities and dynamic systems are standard frameworks for reasoning with time but are not always good modeling languages. This has led in recent years to the development of alternative languages, more suitable for modeling, that can be thought of as one of two types. *Translation languages*, on the one hand, aim to provide ways for specifying logical, probabilistic and deterministic dynamic systems by means of shorter and more intuitive descriptions (e.g., (Pednault 1989; Dean & Kanazawa 1989; Gelfond & Lifschitz 1993)). *Default languages*, on the other, aim to extend classical logic with the ability to express the *expected* effects of actions and the *expected* evolution of fluents (e.g. (McCarthy 1986)).

In this paper we develop a model for temporal reasoning that is hybrid in the sense that it features a language based on defaults and a semantics based on 'approximate' Markov Processes. More precisely, the user describes the dynamics of the domain of interest in terms of default rules, and the defaults get mapped into a Markov Process with probabilities replaced by order-of-magnitude approximations. This results in a framework that relates logical and probabilistic models and accommodates in a natural way features that have been found problematic in other models like non-determinism, action qualifications, parallel actions, and abduction to both fluents and actions.

*Mailing address from US and Europe: Hector Geffner, Bamco CCS 144-00, P.O.BOX 02-5322, Miami Florida 33102-5322, USA. E-mail: hector@usb.ve.

Dynamic Systems

Dynamic systems can often modeled by means of a transition function f that maps states s_i and inputs u_i into unique successor states $s_{i+1} = f(s_i, u_i)$ (Padulo & Arbib 1974; Dean & Wellman 1991).¹ The language for actions developed by Gelfond and Lifschitz (1993) is a language for specifying systems of this sort by means of rules of the form:

$$A \text{ causes } B \text{ if } C \quad (1)$$

Rules such as these are understood as constraints over the function f that must map states s_i where B holds into states s_{i+1} where C holds when the input u_i is A . Under the assumption that B and C are conjunctions of literals, and that all atoms (except actions) *persist by default*, these rules determine the function f completely.

The semantics of Gelfond's and Lifschitz's language is given in terms of such functions. Roughly, a literal L_i follows from a sequence of actions a_0, a_1, \dots and observations o_1, o_2, \dots when L_i is true in all the state-space trajectories s_0, s_1, \dots , that are compatible with the rules (i.e., $s_{i+1} = f(s_i, a_i)$) and the observations (s_i satisfies o_i).

Gelfond and Lifschitz model is not affected by the difficulties reported by Hanks and McDermott (1986) because the transition function f provides the right semantic structure for interpreting defaults in this context. Persistence defaults — which are present in the model even if they are not encoded explicitly by means of rules — are regarded constraints on the possible transitions from one state to the immediate successor states, *independent of both future and past, and the actual observations*. Other models based on a similar idea are Baker's (1991) and Sandewall's (1991).

Markov Processes

The model above assumes that the dynamics of the system is *deterministic* in the sense that knowledge of the state and the inputs is sufficient to predict the

¹This is for systems that are discrete-time, time-invariant and deterministic; see (Padulo & Arbib 1974).

future with complete certainty. For the cases where this assumption is not good a different class of models has been developed in which the state and the inputs predict the future behavior of the system with some known probability (Howard 1971).

Formally, a state s_i and an input u_i determine now a set of possible successor states s_{i+1} with probabilities $P(s_{i+1}|s_i, u_i)$. Then, under the assumptions that future inputs do not affect past states (the causality principle) and that the future is independent of the past given the present (Markovian assumption), the probability of each state *trajectory* s_0, s_1, \dots, s_N given a sequence of inputs u_0, u_1, \dots, u_{N-1} is given by the equation:

$$P(s_0, \dots, s_N | u_0, \dots, u_{N-1}) = P(s_0) \cdot \prod_{i=0}^{N-1} P(s_{i+1} | s_i, u_i) \quad (2)$$

When a set of observations O is obtained, this probability is multiplied by a normalizing constant if the trajectory satisfies the observations, and by zero otherwise. The probability of a proposition is simply the sum of the probabilities of the *trajectories* that make the proposition true.

Models of this type, known as Markov Processes, are significantly more expressive than deterministic models in which transition probabilities can only be zero or one. This generality comes at the price of *specifying* and *computing* with such models. For this reason, attempts to use probabilistic dynamic models in AI have focused on the development of languages that trade off some of the expressive power of probabilistic models for the benefit of simple rule-based specifications (Dean & Kanazawa 1989; Hanks & McDermott 1994).

In this paper we develop a different type of probabilistic temporal model based on rules that may be adequate when exact probabilities are not needed and the distinction between likely and unlike consequences suffices. The key concepts are two: an abstraction of Markov Processes in which probabilities are replaced by their *order-of-magnitude approximations* and a way to specify systems of that sort by means of *partial and incomplete sets of rules*. We consider each idea in turn.

Qualitative Markov Processes

The order-of-magnitude of a probability measure p relative to a small parameter ϵ can be defined as the smallest *integer* $\kappa(p)$ such that $p \leq \epsilon^{\kappa(p)}$. For example, if $\epsilon = 0.2$, the order-of-magnitude of $p_1 = 0.5$ and $p_2 = 0.01$ are $\kappa(p_1) = 0$ and $\kappa(p_2) = 2$ respectively. Interestingly, as shown by Spohn (1988), as the parameter ϵ is made smaller and smaller, in the limit, the order-of-magnitude measures κ obey a calculus given

by the axioms:²

$$\kappa(p) = \min_{w \models p} \kappa(w), \kappa(p \vee \neg p) = 0, \kappa(p|q) = \kappa(p \wedge q) - \kappa(q) \quad (3)$$

which is structurally similar to the calculus of probabilities, with products replaced by sums, sums by minimizations, etc.

Spohn refers to the κ measures as degrees of surprise or disbelief as lower κ measures stand for higher probabilities and higher κ measures stand for lower probabilities. In particular, a proposition p is deemed *plausible* if $\kappa(p) = 0$ and *implausible* or *disbelieved* if $\kappa(p) > 0$. Since the axioms rule out two complementary propositions from being disbelieved at the same time, p is *accepted* or *believed* when its negation is disbelieved, i.e., if $\kappa(\neg p) > 0$.

The appeal of Spohn's κ -calculus is that it combines the basic intuitions underlying probability theory (context dependence, conditionalization, etc.) with the notion of plain belief. The beliefs sanctioned by κ functions are *plain* and *revisable* in the sense that p can be believed given q , and yet $\neg p$ can be believed given q and something else. Indeed, the function κ expresses a preference relations on worlds in which a world w is preferred to w' if $\kappa(w) < \kappa(w')$. From this point of view, the criterion $\kappa(\neg p|q) > 0$ for *accepting* p given q is nothing else but an abbreviation of the standard condition in non-monotonic logics that require p to be true in all maximally preferred worlds that satisfy q (e.g., (Lehmann & Magidor 1988)).

Goldszmidt and Pearl (1992) were the first to exploit the dual connection of the κ -calculus to probabilities and non-monotonic reasoning, showing how some problems in causal default reasoning could be solved by using κ functions that satisfy a *stratification* condition analogous to the condition that defines Probabilistic Bayesian Networks (Pearl 1988). They refer to κ measures as qualitative probabilities, and to stratified κ functions as Qualitative Bayesian Networks (Goldszmidt & Darwiche 1994).

In the same perspective, we consider in this paper *Qualitative Markov Processes*, defined as the κ functions for which the plausibility of trajectories s_0, \dots, s_N given the inputs u_0, \dots, u_{N-1} is given by the qualitative version of (2):

$$\kappa(s_0, \dots, s_N | u_0, \dots, u_{N-1}) = \kappa(s_0) + \sum_{i=0}^{N-1} \kappa(s_{i+1} | s_i, u_i) \quad (4)$$

The model for temporal reasoning below is a language for specifying systems of this sort by means of partial and incomplete sets of rules.

²We also assume $\kappa(p) = \infty$ and $\kappa(q|p) = 0$ when p is unsatisfiable.

The Proposed Model

Language

We deal with temporal models or *theories* specified by means of three type of constructs: *temporal rules*, *observations*, and what we call *completion functions*. The first two are familiar and their precise syntax will be given below. The third one is less familiar and will be introduced in the next section. Intuitively, rules will be restrictions on the possible state transitions, observations will be restrictions on the actual trajectories, and completion functions will be functions that fill out missing information to determine the plausibility of state transitions uniquely.

The syntax of temporal rules and observations presumes a finite set P of time-dependent primitive propositions (atoms) p, q, r, \dots , and a time set T given by the non-negative integers $0, 1, \dots$. We will refer to the language comprised of the propositions in P closed under the standard propositional connectives as \mathcal{L} , and use the symbols A, B , etc. to denote arbitrary formulas in \mathcal{L} . We will also use the symbol L to refer to literals in \mathcal{L} (atoms or their negations) and $\sim L$ to refer to their complements.

The temporal rules are *default* rules of the form $A \rightarrow L$ saying that if A is true at time i , then by default L will be true at time $i+1$ for any $i \in T$. Each rule has a *priority* which is represented by a non-negative integer: the higher the number, the higher the priority. The idea is that when two rules say different things about the same literal, the higher priority rule prevails.

Non-deterministic rules (Sandewal 1991) are accommodated by expressions of the form $A \rightarrow p | \neg p$ and are understood as a *shorthand* for the pair of rules $A \rightarrow p$ and $A \rightarrow \neg p$. Non-deterministic rules express that A sometimes makes p true and sometimes makes p false (e.g., *dropped-cup* \rightarrow *breaks* | \neg *breaks*). Unless otherwise specified, the priority of rules is assumed to be zero (lowest priority).

Finally, the *observations* are formulas that have been observed to be true at some specific time points. We use the notation $p(i)$ to express that the primitive proposition p is true at time i and call such expressions *temporal propositions*. We refer to the language that results from closing such temporal propositions under the standard propositional connectives as \mathcal{L}_T . \mathcal{L}_T will thus be the language of the observations and the conclusions that we may want to draw from them. We will call the formulas in \mathcal{L}_T the *temporal formulas*. Thus, if a, b and c are primitive propositions $a(2) \vee a(3) \supset c(4) \vee \neg b(1)$ will be a valid temporal formula and hence a possible conclusion or observation.

Semantics

A set of temporal rules augmented with a *completion function* will determine one specific Qualitative Markov Process represented by a particular κ function. A conclusion C will then follow from a set of

observations O if $\kappa(\neg C|O) > 0$. We make this precise by defining the states, trajectories, and the form of the κ function.

The *states* are truth assignments to the primitive propositions that determine the truth value of all formulas in \mathcal{L} . The notation s_i, s_j , etc. will be used to denote states at times i, j , etc., while state-space *trajectories* will be denoted as $s_0, s_1, \dots, s_i, \dots$. We will say that a trajectory *satisfies* a temporal proposition $p(i)$ if the state s_i in the trajectory satisfies the primitive proposition p . Following the standard interpretation of the propositional connectives, trajectories represent logical interpretations over the temporal language \mathcal{L}_T assigning a truth-value to all temporal formulas, and hence, to all observations.

For example, a trajectory s_0, s_1, \dots with a state s_1 that satisfies p and q , and a state s_2 that satisfies only p , will satisfy the *temporal* formulas $p(1) \supset p(2)$ and $\neg q(2)$, but not $p(1) \supset q(2)$ or $\neg q(1) \vee \neg p(1)$.

Our dynamic systems will have *no inputs*. Actions, which in other frameworks are represented as inputs, will be represented here in the *state*. Thus to express that a switch was toggled at time $i = 5$, we will simply say that *toggle(5)* was observed. With no inputs, the transition plausibilities that characterize our Markov Processes will simplify from $\kappa(s_{i+1}|s_i, u_i)$ to $\kappa(s_{i+1}|s_i)$. Later on we will show that by representing inputs as observations we are not giving up the ‘causality property’ by which actions should not affect past states (Padulo & Arbib 1974). Instead we will gain the ability to deal naturally with parallel actions, action qualifications and abduction to actions.

Transition Plausibilities. We are left to determine the prior and transition plausibilities $\kappa(s_i)$ and $\kappa(s_{i+1}|s_i)$ from the information provided by the user. Let us say that a state s makes a rule $A \rightarrow L$ *active* in the theory when s satisfies A but s does not satisfy B for any conflicting rule $B \rightarrow \sim L$ with higher priority. Let us also use the notation L_i to stand for temporal literals like $p(i)$ or $\neg p(i)$ and say that L_{i+1} is *supported* by a state s_i when a rule $A \rightarrow L$ is active in s_i . Then the proposed mapping from rules to transition plausibilities $\kappa(s_{i+1}|s_i)$ can be understood in terms of the following assumptions:

1. Literals L_{i+1} and L'_{i+1} that are logically independent are conditionally *independent* given the past s_i .³
2. L_{i+1} is *not disbelieved* when s_i supports L_{i+1}
3. L_{i+1} is *completely disbelieved* when s_i supports $\sim L_{i+1}$ but not L_{i+1}
4. The plausibilities of two complementary literals L_{i+1} and $\sim L_{i+1}$ that are not supported by s_i are *independent* of s_i

³Two literals are logically independent if the truth of one does not constraint the truth of the other.

Assumption 1 excludes the possibility of *ramifications* and translates into the identity:

$$\kappa(s_{i+1}|s_i) = \sum_{L \in s_{i+1}} \kappa(L_{i+1}|s_i) \quad (5)$$

where L ranges over the literals that are true in s_{i+1} . Assumptions 2 and 3 are consequences of the default reading of the rules.⁴ The last assumption is the most important and follows from assuming that *the past influences the future only through the active rules*; i.e., *same active rules mean same transition plausibilities*.

These assumptions impose restrictions on the type of Qualitative Markov Processes that can be expressed, yet with the exception of Assumption 1, we have found them reasonable in domains where predictions can be explained qualitatively in terms of rules and prior judgements. Later on we will show that received models for temporal reasoning that do not deal with ramifications embed these and other assumptions.

The assumptions determine the following mapping from rules to plausibilities:

$$\kappa(L_{i+1}|s_i) = \begin{cases} 0 & \text{when } s_i \text{ supports } L_i \\ \infty & \text{when } s_i \text{ supports } \sim L_i \text{ but not } L_i \\ \pi(L) & \text{when } s_i \text{ supports neither} \end{cases} \quad (6)$$

We call this model of interaction the *Noisy-Rule model* in analogy to the Noisy-OR model used in Bayesian Networks (Pearl 1988) In this model, the parameters $\pi(L)$ and $\pi(\sim L)$ determine the plausibilities of the literals L and $\sim L$ in the absence of reasons to believe in either one of them (see (Geffner 1996) for a different application of this model). The function π is what we call the *completion function* and must be such that for each literal L , $\pi(L)$ must be a non-negative integer and either $\pi(L)$ or $\pi(\sim L)$ must be zero (i.e., π is a plausibility function over L and $\sim L$). For literals L_{i+1} and $\sim L_{i+1}$ that are not supported by *any* state s_i , e.g., the literals L_0 referring to the initial state, $\pi(L_{i+1})$ and $\pi(\sim L_{i+1})$ encode the *prior plausibilities* $\kappa(L_{i+1})$ and $\kappa(\sim L_{i+1})$ respectively.⁵

Two completion functions that we will find useful are the *grounded* and *uniform* functions. The first makes $\pi(\neg p) = 0$ and $\pi(p) = 1$ for each primitive proposition p , expressing that in the absence of reasons for or against p , $\neg p$ is assumed more plausible (like in the Closed World Assumption). The second makes $\pi(p) = \pi(\neg p) = 0$, expressing that in the absence of reasons for or against p , the literals p and $\neg p$ are assumed equally plausible. Later on we will show that some familiar systems embed assumptions that fit naturally with these functions.

⁴We could make Assumption 3 less extreme by replacing 'complete disbelief' (infinite κ) by 'partial disbelief' (non-zero κ). Yet this weaker condition would make the specification of the resulting models more complex.

⁵This is because in that case $\kappa(L_{i+1}) = \min_{s_i} \kappa(L_{i+1}|s_i)$ and $\kappa(L_{i+1}|s_i) = \pi(L)$.

As an illustration, given the rules $q \rightarrow p$ and $r \rightarrow p$, the grounded completion function produces a Noisy-OR type of model in which p is certain given q or r and $\neg p$ is more likely than p when q and r are false (below s_i^+ and s_i^- stand for the states that make $q \vee r$ true or false respectively).

$$\begin{aligned} \kappa(p_{i+1}|s_i^+) &= 0 & \kappa(\neg p_{i+1}|s_i^+) &= \infty \\ \kappa(p_{i+1}|s_i^-) &= 1 & \kappa(\neg p_{i+1}|s_i^-) &= 0 \end{aligned}$$

Summary

The proposed model works as follows. The user provides the rules and the completion function and from (6) we get the plausibilities $\kappa(L_0)$ and $\kappa(L_{i+1}|s_i)$ for all literals and states. These plausibilities are combined by means of (5) to yield the prior and transition plausibilities $\kappa(s_0)$ and $\kappa(s_{i+1}|s_i)$, which plugged into (4) give us the plausibility of any trajectory, and hence, of any formula (3). To determine whether a temporal formula C follows from the observations O we check then whether $\kappa(\neg C|O) > 0$, where $\kappa(\neg C|O)$ is the difference between the plausibilities of the most plausibility trajectories that satisfy $\neg C \wedge O$ and O respectively (3).

Example. Consider the expressions 'if a block is pushed it moves', 'if a block is pushed and is held, it may not move', and 'if a very heavy block is pushed, it does not move', represented by the rules:

$$a \rightarrow p ; a \wedge b \rightarrow p | \neg p ; a \wedge q \rightarrow \neg p$$

in increasing order of priority. We also consider a rule $q \rightarrow q$ capturing the persistence of the property 'heavy', and a grounded completion function π where for every positive literal q , $\pi(\neg q) = 0$ and $\pi(q) = 1$ (i.e., atoms are assumed false by default).

We want to determine whether $p(1)$ follows from $a(0)$; i.e., whether a block will move if pushed. We will use the fact that for the completion function above $\kappa(s_{i+1}|s_i)$, when finite, is equal to the number of atoms x true in s_{i+1} that are not supported by s_i . This also applies to the initial states s_0 where no atom is supported and hence where $\kappa(s_0)$ is equal to the number of atoms true in s_0 .

Consider now the trajectory $t = s_0, s_1, \dots$ that only makes two atoms true: $a(0)$ and $p(1)$. We want to show that t is the single most plausible trajectory compatible with $a(0)$. From the considerations above, $\kappa(s_0) = 1$, and since s_0 supports $p(1)$ but no other atom, $\kappa(s_1|s_0) = 0$. This means that $\kappa(t) = 1$, as all states s_1, s_2, \dots support no atoms and hence $\kappa(s_{i+1}|s_i) = 0$ for all $i > 1$.

We need to show that for any other trajectory $t' = s'_0, s'_1, \dots$ satisfying $a(0)$, $\kappa(t') > 1$. This is actually straightforward as any state s'_0 compatible with $a(0)$ but different than s_0 will have a plausibility $\kappa(s'_0) > 1$, and similarly, any state s'_{i+1} different than s_{i+1} will have a plausibility $\kappa(s'_{i+1}|s_i) > 0$. Thus, t is the *single* most plausible trajectory compatible with $a(0)$, and hence, $a(0)$ implies $p(1)$.

This scenario provides an example of a projection. Examples of parallel actions, non-determinism, action qualifications and abduction can all be obtained by using similar arguments in slightly different settings. For instance, if both $a(0)$ and $b(0)$ are observed (i.e., the block is pushed while held), neither $p(1)$ nor $\neg p(1)$ will be predicted (as both literals are supported by the states s_0 that make $a(0) \wedge b(0)$ true and $q(0)$ false). Similarly, if the observation $q(5)$ is added (the block is very heavy), $\neg p(1)$ will be predicted. Finally, if the rule $\neg a \rightarrow \neg p$ is added, $a(0)$ will follow from $p(1)$ ('the blocked moved, therefore it was pushed').

Action Theories

In this section we will specialize the framework laid out above by introducing some common assumptions about actions and fluents that facilitate the specification and processing of temporal theories. These assumptions are: 1) every primitive proposition represents either an action or a fluent, 2) fluents persist by default, 3) actions are exogenous, 4) actions occur with low probability, and 5) changes occur only in the presence of actions (that are not necessarily known).

Formally, this means that 1) every rule will be a persistence rule or an action rule, 2) persistence rules will be of the form $p \rightarrow p$ and $\neg p \rightarrow \neg p$ (actually we assume one such pair of rules for every fluent p), 3) action symbols a do not occur in the head of the rules, 4) actions are unlikely, i.e., $\pi(a) > 0$, and 5) action rules have priority over persistence rules, and actions symbols are not negated in the body of such rules. We will also assume that all observations are *literals*.

Theories of this type are similar to some of the theories considered in the literature (e.g., Gelfond's and Lifschitz's) yet they allow for non-determinism, arbitrary plausibility function over fluents, abduction to actions and fluents, action qualifications, and parallel actions. We will call such theories, *action theories*.

We mention briefly three main properties of action theories. First, in spite of representing actions as part of the state, actions remain independent of past states in compliance with the so-called *causality principle* (Padulo & Arbib 1974)).⁶

Proposition 1 *In action theories, actions are independent of past states, i.e., if a_i denotes the occurrence of a number of actions at time i , $\kappa(s_i | s_{i-1}, a_{i-1}, a_i, a_{i+1}, \dots) = \kappa(s_i | s_{i-1}, a_{i-1})$.*

Second, all uncertainty in action theories is summarized by the prior plausibility of actions and the prior plausibility of fluents:

Proposition 2 *In action theories, the plausibility measure of a trajectory $t = s_0, s_1, \dots$, when finite,*

⁶Actually, since actions are part of the state, actions affect the present state yet they do not affect the present state of *fluents*.

*is given by the sum of the prior plausibilities of the actions that occur in t and the prior plausibilities of the literal fluents that occur in s_0 .*⁷

$$\kappa(t) = \sum_{L \in s_0} \pi(L) + \sum_i \sum_{a \in s_i} \pi(a)$$

The last property we mention is that only the *relative* prior plausibilities of actions and fluents matter when the theories are *predictive* (i.e., when there are no surprises). For such theories, any two completion functions π and π' that order complementary literals in the same way, i.e., $\pi(L) < \pi(\sim L)$ iff $\pi'(L) < \pi'(\sim L)$, will yield the same behavior.

Definition 1 *An action theory is predictive given a set of observations O and a completion function π , if $\kappa(O | O_A, O_0) = 0$, where $O_A \subseteq O$ refers to the observed actions and $O_0 \subseteq O$ refers the observed fluents at time $i = 0$.*

Proposition 3 *The conclusions sanctioned by a predictive action theory are not affected by changes in the completion function that preserve the plausibility ordering of complementary literals.*

This means that in these theories the exact value of $\pi(L)$ is irrelevant as long as $\pi(L) > 0$, because in such case $\pi(\sim L) < \pi(L)$. For this reason, in such cases it is sufficient to determine whether each positive literal is true by default, false by default, or undecided. The grounded and uniform completion functions, for example, make the second and third choices respectively.

Related Models

The *semantics* of the model draws from approaches in the literature that exploit the double connection of Spohn's κ functions to non-monotonic reasoning and probabilities (Goldszmidt & Pearl 1992; Goldszmidt & Darwiche 1994). The latter work in particular deals with temporal reasoning and is based on structures similar to Bayesian Networks (Pearl 1988) in which conditional probabilities $P(\cdot | \cdot)$ are replaced by conditional plausibilities $\kappa(\cdot | \cdot)$ provided by the user.

The *language* of this model, on the other hand, draws from temporal logics like Gelfond's and Lifschitz's (1993) that do not handle uncertainty. We want to show in this section that the proposed model provides a natural generalization of such logics by representing uncertainty explicitly in the form of completion functions. We focus here on Gelfond's and Lifschitz's logic only; equivalent formalizations are discussed in (Kartta 1993).

For simplicity, and without any loss of generality, we consider domain descriptions with actions rules '*A causes B if C*' and initial conditions '*initially L*'

⁷This is because fluent literals in s_0 and actions anywhere are not supported and hence for them $\kappa(x) = \pi(x)$, while for all other fluent literals $\kappa(L_{i+1} | s_i) = 0$ or $\kappa(L_{i+1} | s_i) = \infty$.

only. Given a domain description D , we define T_D as the action theory with rules $A \wedge B \rightarrow C$ and observations L_0 (all action rules have the same priority, and persistence rules for fluents are implicit as in any action theory). The relation between D and T_D is then as follows:

Proposition 4 *Let D be a consistent domain description.⁸ Then a value proposition ‘ L after A_0, \dots, A_n ’ is entailed by D according to Gelfond and Lifschitz iff the literal L_{n+1} follows from T_D and the actions $A_i(i)$, $i = 0, \dots, n$, under the uniform completion function.*

In other words, Gelfond’s and Lifschitz’s model can be understood in this framework in terms of two assumptions: that complementary fluents are equally plausible, and all actions are implausible a priori. The advantage of the model proposed is that these assumptions are explicit and can be modified by a change in the completion function π (e.g., a grounded completion functions for fluents, for example, leads to the behavior characteristic of negation as failure). This actually explains why we can accommodate action qualifications and represent actions in the state. If we only had the uniform completion function we would get a behavior *monotonic* in the set of actions, very much like Gelfond’s and Lifschitz’s model yields a monotonic behavior in the observations.

Summary

We have developed a qualitative model for temporal reasoning that relates logical and probabilistic approaches, and handles non-determinism, actions qualifications, parallel actions and abduction in a natural way. The model is limited in other ways such as in its inability to deal with ramifications. We hope to address this limitation in the future by making the mapping from rules to transition plausibilities sensitive to the domain constraints. We have also developed inference procedures that we plan to include in the full version of this paper.

Acknowledgments. Many of these ideas originated in conversations with Yoav Shoham. I also want to thank Blai Bonet, Nir Friedman, Joe Halpern for useful discussions, and the anonymous AAAI reviewers for useful comments.

References

- Baker, A. 1991. Nonmonotonic reasoning in the framework of the situation calculus. *AIJ* 49:5–23.
- Dean, T., and Kanazawa, K. 1989. A model for reasoning about persistence and causation. *Computational Intelligence* 5:142–150.
- Dean, T., and Wellman, M. 1991. *Planning and Control*. Morgan Kaufmann.

⁸ D is consistent if no pair of rules associated with the same action have antecedents that are jointly satisfiable and consequents that are jointly unsatisfiable.

- Geffner, H. 1996. A formal framework for causal modeling and argumentation. In *Proceedings FAPR’96*. Bonn. Springer Verlag.
- Gelfond, M., and Lifschitz, V. 1993. Representing action and change by logic programs. *J. of Logic Programming* 17:301–322.
- Goldszmidt, M., and Darwiche, A. 1994. Action networks. In *Proceedings Spring AAAI Symposium on Decision-Theoretic Planning*.
- Goldszmidt, M., and Pearl, J. 1992. Rank-based systems. In *Proceedings KR-92*, 661–672. Morgan Kaufmann.
- Hanks, S., and McDermott, D. 1986. Default reasoning, non-monotonic logics, and the frame problem. In *Proceedings AAAI-86*, 328–333.
- Hanks, S., and McDermott, D. 1994. Modeling a dynamic and uncertain world I: symbolic and probabilistic reasoning about change. *AIJ* 66(1):1–55.
- Howard, R. 1971. *Dynamic Probabilistic Systems—Volume I: Markov Models*. New York: Wiley.
- Kartha, G. 1993. Soundness and completeness results for three formalizations of action. In *Proceedings IJCAI-93*, 724–729.
- Lehmann, D., and Magidor, M. 1988. Rational logics and their models. Dept. of Computer Science, Hebrew University, Israel.
- McCarthy, J. 1986. Applications of circumscription to formalizing commonsense knowledge. *Artificial Intelligence* 28:89–116.
- Padulo, L., and Arbib, M. 1974. *System Theory*. Hemisphere Publishing Co.
- Pearl, J. 1988. *Probabilistic Reasoning in Intelligent Systems*. Los Altos, CA.: Morgan Kaufmann.
- Pednault, E. 1989. ADL: Exploring the middle ground between Strips and the situation calculus. In *Proceedings KR-89*, 324–332.
- Sandewal, E. 1991. Features and fluents. Technical Report R-91-29, CS Department, Linköping University, Linköping, Sweden.
- Spohn, W. 1988. A general non-probabilistic theory of inductive reasoning. In *Proceedings 4th Workshop on Uncertainty*, 315–322.