



**INSTITUT FÜR INFORMATIONSSYSTEME
ABTEILUNG WISSENSBASIERTE SYSTEME**

**PROBABILISTIC REASONING ABOUT ACTIONS
IN NONMONOTONIC CAUSAL THEORIES**

Thomas Eiter and Thomas Lukasiewicz

INFSYS RESEARCH REPORT 1843-03-01
MARCH 2003

Institut für Informationssysteme
Abtg. Wissensbasierte Systeme
Technische Universität Wien
Favoritenstraße 9-11
A-1040 Wien, Austria
Tel: +43-1-58801-18405
Fax: +43-1-58801-18493
sek@kr.tuwien.ac.at
www.kr.tuwien.ac.at



TECHNISCHE UNIVERSITÄT WIEN

INFSYS RESEARCH REPORT

INFSYS RESEARCH REPORT 1843-03-01, MARCH 2003

PROBABILISTIC REASONING ABOUT ACTIONS
IN NONMONOTONIC CAUSAL THEORIES

(PRELIMINARY REPORT, 20 JUNE 2003)

Thomas Eiter¹ and Thomas Lukasiewicz²

Abstract. The action language $\mathcal{C}+$ is an important high-level formalism for describing actions, which has evolved from a number of earlier proposals for action languages, and which is based on McCain and Turner's nonmonotonic causal logics. In this paper, we present the action language $\mathcal{PC}+$ for probabilistic reasoning about actions, which is a generalization of $\mathcal{C}+$ that allows to deal with probabilistic as well as nondeterministic effects of actions. We define a formal semantics of $\mathcal{PC}+$ in terms of probabilistic transitions between sets of states. Using the concept of a history and its belief state, we then show how several important problems in reasoning about actions can be concisely formulated in our formalism. A history h is a sequence of actions and observations, which are labeled with a reasoning modality over sets of states. It has an associated belief state, which comprises possible sets of states and probabilistic information to describe the qualitative and quantitative knowledge about h . As for computation, we present a compact representation of belief states and prove its correctness in implementing belief states. We also show how it can be used to reduce reasoning about actions in $\mathcal{PC}+$ to reasoning in nonmonotonic causal theories, which is a step towards implementation on top of existing technology for nonmonotonic causal theories such as the Causal Calculator.

¹Institut für Informationssysteme, Technische Universität Wien, Favoritenstraße 9-11, A-1040 Vienna, Austria; e-mail: eiter@kr.tuwien.ac.at.

²Dipartimento di Informatica e Sistemistica, Università di Roma "La Sapienza", Via Salaria 113, I-00198 Rome, Italy; e-mail: lukasiewicz@dis.uniroma1.it. Alternate address: Institut für Informationssysteme, Technische Universität Wien, Favoritenstraße 9-11, A-1040 Vienna, Austria; e-mail: lukasiewicz@kr.tuwien.ac.at.

Acknowledgements: This work has been partially supported by FWF (Austrian Science Funds) under the projects P14781-INF, P16536-N04, and Z29-N04 and by the European Commission under the grant FET-2001-37004 WASP and the Marie Curie Individual Fellowship HPMF-CT-2001-001286 of the programme "Human Potential" (Disclaimer: The authors are solely responsible for information communicated and the European Commission is not responsible for any views or results expressed).

Copyright © 2003 by the authors

Contents

1	Introduction	1
2	Syntax and Semantics of $\mathcal{PC}+$	3
2.1	Preliminaries	3
2.2	Syntax of $\mathcal{PC}+$	3
2.3	Semantics of $\mathcal{PC}+$	5
3	Histories and Belief States for $\mathcal{PC}+$	8
3.1	Belief States	8
3.2	Compact Belief States	9
4	Reasoning and Planning in $\mathcal{PC}+$	10
4.1	Prediction	10
4.2	Postdiction	12
4.3	Planning	13
5	Conditional Planning in $\mathcal{PC}+$	13
6	Reduction to Causal Theories	15
7	Summary and Outlook	16
A	Proofs	16

1 Introduction

One of the most crucial problems that we have to face in reasoning about actions for mobile robotics in real-world environments is uncertainty, both about the initial situation of the robot's world and about the results of the actions taken by the robot (due to noisy effectors and/or sensors). One way of adding uncertainty to reasoning about actions is based on qualitative models in which all possible alternatives are equally taken into consideration. Another way is based on quantitative models where we have a probability distribution on the set of possible alternatives, and thus can numerically distinguish between possible alternatives.

Well-known first-order formalisms for reasoning about actions such as the situation calculus [19] easily allow for expressing qualitative uncertainty about the effects of actions and the initial situation of the world through disjunctive knowledge. Furthermore, there are generalizations of the action language \mathcal{A} [6] that allow for qualitative uncertainty in the form of nondeterministic actions. An important recent formalism in this family is the action language $\mathcal{C}+$ [7], which is based on the theory of nonmonotonic causal reasoning presented in [13], and has evolved from the action language \mathcal{C} [8]. In addition to allowing for conditional and nondeterministic effects of actions, $\mathcal{C}+$ also supports concurrent actions [2] as well as indirect effects and preconditions of actions through static causal laws [20]. Closely related to it is the recent planning language \mathcal{K} [4].

There are a number of formalisms for probabilistic reasoning about actions. In particular, Bacchus et al. [1] propose a probabilistic generalization of the situation calculus, which is based on first-order logics of probability, and which allows to reason about an agent's probabilistic degrees of belief and how these beliefs change when actions are executed. Poole's independent choice logic [16, 17, 18] is based on acyclic logic programs under different "choices". Each choice along with the acyclic logic program produces a first-order model. By placing a probability distribution over the different choices, we then obtain a distribution over the set of first-order models. Other probabilistic extensions of the situation calculus are given in [12, 5]. A probabilistic extension of the action language \mathcal{A} is given in [3].

The main idea behind the present paper is to orthogonally combine qualitative and quantitative uncertainty in a uniform framework for reasoning about actions: Even though there is extensive work on qualitative and quantitative models separately, there is only few work on such combinations. One such approach is due to Halpern and Tuttle [10], which combines nondeterminism and probabilistic uncertainty in a game-theoretic framework. Halpern and Tuttle argue in particular that "some choices in a distributed system must be viewed as inherently nondeterministic (or, perhaps better, nonprobabilistic), and that it is inappropriate, both philosophically and pragmatically, to model probabilistically what is inherently nondeterministic". This underlines the strong need for explicitly modeling qualitative uncertainty in addition to probabilistic uncertainty.

In this paper, we combine the qualitative uncertainty in the action language $\mathcal{C}+$ with probabilistic uncertainty as in [16, 17, 18]. The main contributions of this paper are summarized as follows:

- We present the language $\mathcal{PC}+$ for probabilistic reasoning about actions, which is a probabilistic generalization of the action language $\mathcal{C}+$. It allows for representing actions with conditional and indirect effects, nondeterministic actions, and concurrently executed actions as the main features of $\mathcal{C}+$ as well as probabilistic knowledge about the effects of actions and the initial situation of the world.
- As a central property, $\mathcal{PC}+$ combines in a single framework qualitative as well as quantitative uncertainty, both about the effects of actions and about the initial situation of the world. Here, qualitative uncertainty is represented by forming a set of possible alternatives, while quantitative uncertainty is expressed through a probability distribution on a set of possible sets of possible alternatives.

- We define a formal semantics of $PC+$ by interpreting probabilistic action descriptions in $PC+$ as probabilistic transitions as in partially observable Markov decision processes (POMDPs) [11]. However, it is important to point out that these probabilistic transitions are between *sets of states* rather than *single states*. It is this which allows to handle qualitative uncertainty in addition to the quantitative probabilistic uncertainty as in POMDPs. Differently from standard POMDPs, our approach here only allows for observations without noise, but not for noisy sensing. It also does not deal with costs/rewards of actions.
- We define histories and their belief states in $PC+$. Informally, a history h is a sequence of actions and observations, which are labeled with a reasoning modality over sets of states. It has an associated belief state, which comprises possible sets of states and probabilistic information to describe the qualitative and quantitative knowledge about h . Note that such belief states model partial observability. Moreover, we present a compact representation of belief states, which is based on the notion of a context to encode possible sets of states, and prove a representation theorem, which can be exploited in implementation.
- We show how to express a number of important problems in probabilistic reasoning about actions (namely, the problems of prediction, postdiction, and planning; see especially [19, 12, 7]) in terms of belief states in $PC+$. We also show how to formulate conditional planning in $PC+$.
- We discuss how to reduce probabilistic reasoning about actions in $PC+$ to reasoning in nonmonotonic causal theories [7]. This is a step towards implementation on top of existing technology (such as the Causal Calculator [14]).

The work closest in spirit to this paper is perhaps the one by Baral et al. [3], which also proposes a logic-based formalism for probabilistic reasoning about actions. However, there are several crucial differences. First, and as a central conceptual difference, our work orthogonally combines quantitative and qualitative uncertainty, in the form of probability distributions and sets of possible alternatives, respectively, while Baral et al. only allow for quantitative uncertainty in the form of probability distributions. Note that Baral et al. circumvent the problem of dealing with qualitative uncertainty by making the strong assumption of a uniform distribution whenever the probabilities for possible alternatives are not known. Second, Baral et al. allow only for a quite restricted form of probability distributions, which are either uniform distributions or produced from uniform distributions. Third, our language $PC+$ generalizes the action language $\mathcal{C}+$, while Baral et al.'s language generalizes the action language \mathcal{A} . Note that $\mathcal{C}+$ is a novel action language that evolved from \mathcal{A} and that is much more expressive than \mathcal{A} .

Another important formalism that is related to ours is Poole's independent choice logic [16, 17, 18], which uses a similar way of adding probabilities to an approach based on acyclic logic programs. But also here the central conceptual difference is that Poole's independent choice logic does not allow for qualitative uncertainty in addition to quantitative uncertainty. Poole circumvents the problem of dealing with qualitative uncertainty by imposing the strong condition of acyclicity on logic programs. Furthermore, Poole's work is more inspired by the situation calculus and less by the action languages around \mathcal{A} .

The rest of this paper is organized as follows. In Section 2, we define $PC+$ and its semantics in probabilistic transitions between sets of states. Section 3 introduces histories, belief states, and compact belief states. In Section 4, we describe how several important problems in probabilistic reasoning about actions can be expressed in our framework. Section 6 shows how to reduce reasoning in $PC+$ to reasoning in nonmonotonic causal theories. In Section 7, we summarize the main results and give an outlook on future research. Note that detailed proofs of all results are given in Appendix A.

2 Syntax and Semantics of $PC+$

In this section, we first recall nonmonotonic causal theories from [7]. We then present the language $PC+$ for probabilistic reasoning about actions, and give an example of a probabilistic action description and initial database expressed in $PC+$. We finally define the semantics of $PC+$ through probabilistic transitions between sets of states.

Informally, the main idea behind our probabilistic extension of $C+$ is to associate with the initial database and with every stochastic action a probability distribution on a set of *contexts*, which are values of exogenous variables. Every sequence of actions from an initial database is then associated with a probability distribution on a set of combined contexts. Hence, probabilistic reasoning about actions in $PC+$ can essentially be reduced to standard reasoning about actions in $C+$ with respect to such combined contexts. Note that Poole’s independent choice logic [16, 17, 18] uses a similar way of adding probabilities to an approach based on acyclic logic programs.

2.1 Preliminaries

We now recall (nonmonotonic) causal theories from [7], which are used to specify initial sets of states and transitions from states to sets of states through actions. Roughly, a causal theory T is a set of ‘causal rules’ $F \Leftarrow G$ with the meaning ‘if G holds, then there is a cause for F ’. In this paper, we consider only finite T . We now first define the syntax of causal theories and then their semantics.

We assume a finite set of *variables* \mathcal{X} . Every variable $X \in \mathcal{X}$ may take on *values* from a finite *domain* $I(X)$. We define *formulas* by induction as follows. *False* and *true*, denoted \perp and \top , respectively, are formulas. If $X \in \mathcal{X}$ and $x \in I(X)$, then $X = x$ is a formula (called *atom*). If F and G are formulas, then also $\neg F$ and $(F \wedge G)$ are formulas. As usual, $X \neq x$ abbreviates $\neg X = x$. A (*causal*) *rule* is an expression of the form $F \Leftarrow G$, where F and G are formulas, called its *head* and *body*, respectively. A *causal theory* is a finite set of rules.

An *interpretation* I of $\mathcal{Y} \subseteq \mathcal{X}$ maps every $Y \in \mathcal{Y}$ to an element of $I(Y)$. We use $I(\mathcal{Y})$ to denote the set of all interpretations of \mathcal{Y} . We obtain $I(F)$ and $I(F \Leftarrow G)$ from F and $F \Leftarrow G$, respectively, by replacing every atom $Y = y$ such that $Y \in \mathcal{Y}$ and $I(Y) = y$ (resp., $I(Y) \neq y$) by \top (resp., \perp). An interpretation I of \mathcal{Y} *satisfies* an atom $Y = y$ with $Y \in \mathcal{Y}$, denoted $I \models Y = y$, iff $I(Y) = y$. Satisfaction is extended to all formulas over \mathcal{Y} as usual (that is, $I \models \neg F$ iff $I \models F$ does not hold, and $I \models (F \wedge G)$ iff $I \models F$ and $I \models G$).

Let T be a causal theory and I be an interpretation of the variables in T . The *reduct* of T relative to I , denoted T^I , is defined as $\{F \Leftarrow G \in T, I \models G\}$. The interpretation I is a *model* of T , denoted $I \models T$, iff I is the unique model of T^I . The theory T is *consistent* iff it has a model.

2.2 Syntax of $PC+$

We next define the syntax of the probabilistic action language $PC+$, which generalizes the syntax of $C+$ [7]. We refer to [7] for further motivation and background for $C+$. We illustrate the language $PC+$ along a (simplistic) robot action domain, which is summarized in Fig. 1. This example shows in particular how quantitative as well as qualitative uncertainty about both the effects of actions and the initial situation of the world can be encoded in $PC+$.

We divide the variables in \mathcal{X} into *rigid*, *fluent*, *action*, and *context* variables. The fluent variables (or *fluents*) are additionally divided into *simple* and *statically determined* ones. We assume that action variables have the Boolean domain $\{\perp, \top\}$. Intuitively, the world is described through rigid variables and fluents. The values of rigid variables do not change when actions are performed, while those of simple (resp.,

statically determined) fluents may directly (resp., indirectly) change through actions. Action variables are used to describe actions, while context variables allow for adding probabilistic knowledge about the effects of actions and about the initial situation of the world.

Example 2.1 In the robot action domain in Fig. 1, a mobile robot r may move to the locations a , b , and c , and carry one of two objects o_1 and o_2 after pickup. This world is described through the simple fluents $at(o_1)$, $at(o_2)$, and $at(r)$ with the domain $\{a, b, c, lost\}$, where $at(O) = L$ iff O is at location L . Moreover, we have the simple fluent $holds$ with the domain $\{o_1, o_2, nil\}$, where $holds = O$ iff r holds O . We then have the action variables $goto(a)$, $goto(b)$, $goto(c)$, $pickup$, and $drop$, which represent the elementary actions “move to location a ”, “move to location b ”, “move to location c ”, “pick up an object”, and “drop an object”, respectively. Finally, the action “move to location L ”, where $L \in \{a, b, c\}$, succeeds only with a certain probability. To model this, we use the context variables $c_g(a)$, $c_g(b)$, and $c_g(c)$ with the domain $\{ok, fail\}$, where $c_g(L) = ok$ iff “move to location L ” succeeds, for every $L \in \{a, b, c\}$. \square

We next define static causal laws, which represent knowledge about fluents and rigid variables. Formally, a *static causal law* is an expression of the form

$$\mathbf{caused} F \text{ if } G, \tag{1}$$

where F and G are formulas such that either (a) every variable occurring in F is a fluent, and no variable occurring in G is an action variable, or (b) every variable occurring in F or G is rigid. If $G = \top$, then (1) is abbreviated by **caused** F . Roughly, (1) encodes that every state of the world that satisfies G should also satisfy F . More formally, (1) is interpreted as the causal rule $F \Leftarrow G$.

Example 2.2 The static causal law (6) **caused** $at(O) = L$ **if** $holds = O \wedge at(r) = L$ expresses that if the robot r is at location L , and r holds O , then O is also at location L . \square

We now define dynamic causal laws, which express how the simple fluents change through actions, and which also encode execution denials for actions. Formally, a *dynamic causal law* is an expression

$$\mathbf{caused} F \text{ if } G \text{ after } H, \tag{2}$$

where F , G , and H are formulas such that (i) every variable occurring in F is a simple fluent, (ii) no variable in G is an action variable, and (iii) no variable in H is a context variable. If $G = \top$, then (2) is abbreviated by **caused** F **after** H . We use **inertial** X to abbreviate the set of all rules (2) such that $F = G = H = X = x$ and $x \in I(X)$. If $F = \perp$ and $G = \top$, then (2) is called an *execution denial* and abbreviated by

$$\mathbf{nonexecutable} H. \tag{3}$$

Roughly, (2) expresses that every possible next state of the world that satisfies G should also satisfy F , if the current state and the executed action satisfy H . More formally, (2) is interpreted as the causal rule $F \Leftarrow G \wedge H$, where G and F refer to the possible next states of the world, and H refers to the current state and the executed action.

Example 2.3 The dynamic causal law (13) **caused** $holds = nil$ **after** $drop$ says that r holds nothing after $drop$. The execution denial (12) expresses that $pickup$ cannot be executed if r already holds an object or if there is no object at the same location as r . The dynamic causal law (11) says that r cannot pick up an object that is not at the same location as r , and (10) says that r holds o_1 respectively o_2 after $pickup$. Thus, there is qualitative uncertainty in the effects of $pickup$: if both o_1 and o_2 are at the same location as r , then $pickup$ results in r picking up either o_1 or o_2 , but it is unpredictable which object r actually picks up. \square

A *causal law* (or *axiom*) is a static or dynamic causal law. Our causal laws generalize their classical counterparts from [7] in the sense that they may also contain context variables. We next introduce the new concept of a context law. Context variables along with such context laws allow for expressing probabilistic effects of actions and probabilistic knowledge about the initial situation of the world.

More formally, a *dynamic context law* for a context variable $X \in \mathcal{X}$ is an expression of the form

$$X = (x_1 : p_1, \dots, x_n : p_n) \text{ after } A, \quad (4)$$

where (i) $I(X) = \{x_1, \dots, x_n\}$, (ii) $p_1, \dots, p_n > 0$, (iii) $p_1 + \dots + p_n = 1$, and (iv) A is a formula over action variables. We use $Pr(X = x_i)$ to denote p_i . If $A = \top$, then (4) is called a *static context law* and abbreviated by

$$X = (x_1 : p_1, \dots, x_n : p_n). \quad (5)$$

Roughly, (4) encodes that after executing an action that satisfies A , the probability that X has the value x_i is given by p_i . Note that a possible generalization of context laws could be to specify a set of probability distributions rather than a single probability distribution.

Example 2.4 The actions $goto(a)$, $goto(b)$, and $goto(c)$ succeed only with certain probabilities. This is modeled using the context variables $c_g(a)$, $c_g(b)$, and $c_g(c)$ in the dynamic causal laws (7) and (8), along with the dynamic context laws (14)–(16). For example, the probability that r really arrives at a after executing $goto(a)$ is given by 0.95. \square

We next define the notion of a probabilistic action description (resp., probabilistic initial database), which encodes the effects of all actions (resp., the initial situation of the world).

Definition 2.5 A *probabilistic action description* D is a finite set of causal and dynamic context laws such that D contains exactly one dynamic context law for every context variable $X \in \mathcal{X}$ in D . Any such D is a *probabilistic initial database*, if all causal and context laws in D are static.

Example 2.6 In the robot action domain in Fig. 1, the probabilistic action description D is completely given by the sentences (6)–(17). Here, axioms (6) and (17) take care of the well-known ramification and frame problem, respectively. Note that axiom (18) forbids concurrent actions.

A probabilistic initial database D_0 may be given as follows. The initial locations of o_1 and o_2 are known with probabilities, which we express by the context variables $c_{at(o_1)}$ and $c_{at(o_2)}$, the static context laws $c_{at(o_1)} = (a : 0.1, b : 0.8, c : 0.1)$ and $c_{at(o_2)} = (a : 0.3, b : 0.6, c : 0.1)$, respectively, and the static causal laws **caused** $at(O) = L$ if $c_{at(O)} = L$, where $O \in \{o_1, o_2\}$ and $L \in \{a, b, c\}$. For example, object o_1 is at location a with a probability of 0.1. Moreover, r is at a or b , expressed by **caused** $at(r) = a \vee at(r) = b$. Finally, r holds no object, expressed by **caused** $holds = nil$. \square

In the sequel, D (resp., D_0) denotes a probabilistic action description (resp., probabilistic initial database).

2.3 Semantics of PC+

We now define the semantics of PC+. Informally, certain interpretations of rigid and fluent variables serve as possible states of the world. We then associate with D_0 a collection of sets of such states, where each set of states has an associated probability. Furthermore, we associate with D a mapping that assigns to

-
- (i) simple fluents: $at(O)$, $O \in \{o_1, o_2, r\}$: $\{a, b, c, lost\}$;
 $holds$: $\{o_1, o_2, nil\}$.
- (ii) action variables: $goto(L)$, $L \in \{a, b, c\}$; $pickup$; $drop$.
- (iii) context variables: $c_g(L)$, $L \in \{a, b, c\}$: $\{ok, fail\}$.
- (iv) static causal laws:
caused $at(O) = L$ **if** $holds = O \wedge at(r) = L$ (6)
- (v) dynamic causal laws:
caused $at(r) = L$ **if** $c_g(L) = ok$ **after** $goto(L)$ (7)
caused $at(r) = L'$ **if** $at(r) = L' \wedge c_g(L) = fail$ (8)
after $goto(L)$ (for $L' \neq L$)
nonexecutable $goto(L) \wedge at(r) = L$ (9)
caused $holds = O$ **if** $holds = O$ (10)
after $pickup$ (for $O \in \{o_1, o_2\}$)
caused \perp **if** $holds = O$ (11)
after $pickup \wedge at(r) = L \wedge at(O) \neq L$
nonexecutable $pickup \wedge [holds \neq nil$ (12)
 $\vee (at(r) = L \wedge at(o_1) \neq L \wedge at(o_2) \neq L)]$
caused $holds = nil$ **after** $drop$ (13)
- (vi) dynamic context laws:
 $c_g(a) = (ok : 0.95, fail : 0.05)$ **after** $goto(a)$ (14)
 $c_g(b) = (ok : 0.95, fail : 0.05)$ **after** $goto(b)$ (15)
 $c_g(c) = (ok : 0.90, fail : 0.10)$ **after** $goto(c)$ (16)
- (vii) inertial laws: for all simple fluents f :
inertial f (17)
- (viii) other execution denials: for all action variables $a_1 \neq a_2$:
nonexecutable $a_1 \wedge a_2$ (18)
-

Figure 1: Robot Action Domain

each pair (S, σ) , consisting of a current set of states S and a labeled action or observation σ , a probabilistic distribution on a collection of future sets of states. Thus, we interpret D by probabilistic transitions as in partially observable Markov decision processes (POMDPs) [11]. But the probabilistic transitions here are between *sets of states* rather than *single states*, which allows to handle qualitative uncertainty in addition to quantitative probabilistic uncertainty. Actions and observations are treated in a uniform way and labeled with modalities to specify how their preconditions (resp., observed formulas) are evaluated on sets of states (see also Section 4).

Semantics of D_0 . Informally, we associate with D_0 a finite set of *contexts* γ , where each γ is in turn associated with a probability value $Pr_0(\gamma)$ and a set of states Φ_γ . Thus, D_0 is interpreted as the collection of all Φ_γ , where each Φ_γ has the probability $Pr_0(\gamma)$. We say that D_0 is *consistent* iff $\Phi_\gamma \neq \emptyset$ for all contexts γ .

Formally, let \mathcal{X}_0 denote the set of all context variables in D_0 . We call $\gamma \in I(\mathcal{X}_0)$ a *context* for D_0 . Its *probability*, denoted $Pr_0(\gamma)$, is defined as $\prod_{X \in \mathcal{X}_0} Pr(X = \gamma(X))$. For $\mathcal{X}_0 = \emptyset$, the empty mapping $\gamma = \emptyset$ is the only context for D_0 , which has the probability $Pr_0(\gamma) = 1$. For each $\gamma \in I(\mathcal{X}_0)$, we define Φ_γ as the set of all models over rigid and fluent variables of the causal theory comprising all $\gamma(F \Leftarrow G)$ for each axiom (1) in D_0 and all $X = x \Leftarrow X = x$ for each simple fluent $X \in \mathcal{X}$ and $x \in I(X)$.

To define the semantics of D , we now formally define states, actions, and observations. We also define how a context γ , current state s , and action or observation σ is associated with a set of future states $\Phi_\gamma(s, \sigma)$.

States. A *state* s is either a member of some Φ_γ , or a model over rigid and fluent variables of the causal theory comprising all $\gamma(F \Leftarrow G)$ for each axiom (1) in D , for any interpretation γ of the context variables in \mathcal{X} , and all $X = x \Leftarrow X = x$ for each simple fluent $X \in \mathcal{X}$ and $x \in I(X)$.

Actions. An *action* α is an interpretation of the action variables in \mathcal{X} . Intuitively, each action variable is a basic action, and α is the concurrent execution of all basic actions that are true under α . The *precondition* for α , denoted π_α , is the conjunction of all $\neg H$ for every execution denial (3) in D such that $s \cup \alpha \models H$ for some state s . An action α is *executable* in a state s , denoted $\pi_\alpha(s)$, iff $s \cup \alpha \models \pi_\alpha$.

We next associate with every action α a set of contexts, and with each such context γ a probability $Pr_\alpha(\gamma)$ and a mapping from states s to a set of future states $\Phi_\gamma(s, \alpha)$. Intuitively, if α is executed in the state s under the context γ , then the set of future states is given by $\Phi_\gamma(s, \alpha)$.

Formally, for states s such that $\pi_\alpha(s)$, denote by $\mathcal{X}_{s, \alpha}$ the set of all context variables in some axiom (1) or (2) in D such that $s \cup \alpha \models H$. We define \mathcal{X}_α as the union of all $\mathcal{X}_{s, \alpha}$. We call $\gamma \in I(\mathcal{X}_\alpha)$ a *context* for α . Its *probability*, denoted $Pr_\alpha(\gamma)$, is defined as $\prod_{X \in \mathcal{X}_\alpha} Pr(X = \gamma(X))$.

An *action transition* is a triple (s, α, s') , where s and s' are states such that $s(X) = s'(X)$ for every rigid variable $X \in \mathcal{X}$, and α is an action such that $\pi_\alpha(s)$. A formula F is *caused* in (s, α, s') under $\gamma \in I(\mathcal{X}_\alpha)$ iff D contains either (a) an axiom (1) such that $s' \cup \gamma \models G$, or (b) an axiom (2) such that $s \cup \alpha \models H$ and $s' \cup \gamma \models G$. We say (s, α, s') is *causally explained* under γ iff s' is the only interpretation that satisfies all formulas caused in (s, α, s') under γ . For every state s and action α , we define $\Phi_\gamma(s, \alpha)$ as the set of all s' such that (s, α, s') is *causally explained* under γ . Note that $\Phi_\gamma(s, \alpha) = \emptyset$ if no such (s, α, s') exists, in particular, if the action α is not executable in the state s .

Observations. An *observation* ω is a formula over fluents. For states s , we use $\pi_\omega(s)$ to abbreviate $s \models \omega$. To treat actions and observations in a uniform way, we also associate with every observation ω a set of contexts γ , and with each γ a probability $Pr_\omega(\gamma)$ and a mapping from states s to a set of future states $\Phi_\gamma(s, \omega)$.

We define $\mathcal{X}_\omega = \emptyset$. The empty mapping $\gamma \in I(\mathcal{X}_\omega)$ is the *context* for ω . It has the *probability* $Pr_\omega(\gamma) = 1$. For states s and observations ω , we define $\Phi_\gamma(s, \omega) = \{s\}$, if $\pi_\omega(s)$, and $\Phi_\gamma(s, \omega) = \emptyset$, otherwise.

We are now ready to define the semantics of D .

Semantics of D . Intuitively, we use D to associate with sets of states S , and actions or observations σ a probability distribution on future sets of states $Pr_\sigma(\cdot | S)$. We say D is *consistent* iff $\Phi_\gamma(s, \alpha) \neq \emptyset$ for all states s , actions α such that $\pi_\alpha(s)$, and contexts $\gamma \in I(\mathcal{X}_\alpha)$. We now first extend $\pi_\sigma(s)$ and $\Phi_\gamma(s, \sigma)$ from states s to sets of states S .

To specify how preconditions of actions (resp., observed formulas) are evaluated on sets of states, we add modality labels to actions (resp., observations). Formally, a *labeled action* (resp., *labeled observation*) is of the form $\circ\tau$, where $\circ \in \{\diamond, \square\}$ and τ is an action (resp., observation). For sets of states S , we use $\pi_{\circ\tau}(S)$ (resp., $\pi_{\square\tau}(S)$) to denote $\exists s \in S: \pi_\tau(s)$ (resp., $\forall s \in S: \pi_\tau(s)$). We define $\mathcal{X}_{\circ\tau} = \mathcal{X}_\tau$ and $Pr_{\circ\tau} = Pr_\tau$.

For every set of states S , every labeled action or observation $\sigma = \circ\tau$ with $\pi_\sigma(S)$, and every context $\gamma \in I(\mathcal{X}_\sigma)$, we then define $\Phi_\gamma(S, \sigma) = \bigcup_{s \in S} \Phi_\gamma(s, \tau)$. Observe that for observations ω , it then holds that $\Phi_\gamma(S, \diamond\omega) = \{s \in S \mid s \models \omega\}$ and $\Phi_\gamma(S, \square\omega) = S$.

We now define the probabilistic transition between sets of states S and S' under σ with $\pi_\sigma(S)$ by:

$$Pr_\sigma(S' | S) = \sum_{\gamma \in I(\mathcal{X}_\sigma), S' = \Phi_\gamma(S, \sigma)} Pr_\sigma(\gamma).$$

Intuitively, given any set of states S such that $\pi_\sigma(S)$, the $\Phi_\gamma(S, \sigma)$'s are the future sets of states under σ , where each $\Phi_\gamma(S, \sigma)$ has the probability $Pr_\sigma(\gamma)$.

Assumption 2.7 In the rest of this paper, we implicitly assume that D and D_0 are consistent, and that all static causal laws (1) in D over rigid variables also belong to D_0 .

3 Histories and Belief States for PC+

Our framework for reasoning and planning in PC+ involves finite sequences of labeled actions and observations, called *histories*, which are inductively defined as follows. The *empty history* ε is a history. If h is a history, and σ is a labeled action or observation, then h, σ is a history. Histories ε, r are abbreviated by r . The *action length* of a history h , denoted $alen(h)$, is the number of occurrences of actions in h .

Example 3.1 In the running example, $\diamond goto(b), \square pickup, \square goto(c), \square at(o_1)=c \vee at(o_2)=c$ is a history of action length 3. Informally, it represents the statement ‘if $goto(b)$ has been executed, then $pickup, goto(c)$ can be executed, and $at(o_1)=c \vee at(o_2)=c$ is observed after that’. \square

We use the notion of a belief state to describe the probabilistic information associated with a history h . Intuitively, a belief state consists of a probability value for h and a probability function on a set of state sets. We now first define belief states and then a compact representation of them.

3.1 Belief States

The *belief state* $b_h = (p_h, \mathcal{S}_h, Pr_h)$ for a history h consists of a real number $p_h \in [0, 1]$ (called *probability* of h , denoted $Pr(h)$), a set of state sets \mathcal{S}_h , and a probability function Pr_h on \mathcal{S}_h . It is inductively defined by:

- If $h = \varepsilon$, then $p_h = 1$, $\mathcal{S}_h = \{\Phi_\gamma \mid \gamma \in I(\mathcal{X}_0)\}$, and for all $S \in \mathcal{S}_h$, $Pr_h(S) = \sum_{\gamma \in I(\mathcal{X}_0), S = \Phi_\gamma} Pr_0(\gamma)$.

- If $h = r, \sigma$, then b_h is defined iff (i) $b_r = (p_r, \mathcal{S}_r, Pr_r)$ is defined, and (ii) $\pi_\sigma(S)$ for some $S \in \mathcal{S}_r$. If b_h is defined, then it is given as follows:

$$\begin{aligned} p_h &= p_r \cdot \sum_{S \in \mathcal{S}_r, \pi_\sigma(S)} Pr_r(S) \\ \mathcal{S}_h &= \{\Phi_\gamma(S, \sigma) \mid \gamma \in I(\mathcal{X}_\sigma), S \in \mathcal{S}_r, \pi_\sigma(S)\} \\ Pr_h(S') &= \frac{p_r}{p_h} \cdot \sum_{\substack{S \in \mathcal{S}_r, \\ \pi_\sigma(S)}} Pr_\sigma(S'|S) \cdot Pr_r(S) \quad (\forall S' \in \mathcal{S}_h). \end{aligned}$$

Intuitively, b_ε describes the probabilistic knowledge associated with D_0 , while $b_{r, \sigma}$ represents the probabilistic knowledge about the world after the history $r, \circ \alpha$ (resp., $r, \circ \omega$), which depends on (i) our respective knowledge after r , and (ii) the effects of α encoded in D (resp., the observation ω).

The following result shows that $\sigma = \circ \omega$ corresponds to a conditioning of Pr_r on all state sets $S \in \mathcal{S}_r$ with $\pi_\sigma(S)$, along with removing from such S all states violating ω .

Proposition 3.2 *Let $h = r, \sigma$ be a history, where $\sigma = \circ \omega$ is a labeled observation. Then, $b_h = (p_h, \mathcal{S}_h, Pr_h)$ is defined iff (i) $b_r = (p_r, \mathcal{S}_r, Pr_r)$ is defined, and (ii) some $S \in \mathcal{S}_r$ exists with $\pi_\sigma(S)$. If b_h is defined, then it is given by:*

$$\begin{aligned} p_h &= p_r \cdot \sum_{S \in \mathcal{S}_r, \pi_\sigma(S)} Pr_r(S) \\ \mathcal{S}_h &= \{\{s \in S \mid s \models \omega\} \mid S \in \mathcal{S}_r, \pi_\sigma(S)\} \\ Pr_h(S') &= \frac{p_r}{p_h} \cdot \sum_{\substack{S \in \mathcal{S}_r, \pi_\sigma(S), \\ S' = \{s \in S \mid s \models \omega\}}} Pr_r(S) \quad (\forall S' \in \mathcal{S}_h). \end{aligned} \tag{19}$$

3.2 Compact Belief States

A sequence of belief states for a sequence of labeled actions and observations can be more compactly represented as a collection of sequences of sets of states, along with a probability for each such sequence. Following this idea, we now introduce the concept of a compact belief state.

For every history h of action length i , we define the set of variables \mathcal{X}_h by induction as follows:

$$\mathcal{X}_h = \begin{cases} 0 : \mathcal{X}_0 & \text{if } h = \varepsilon \\ \mathcal{X}_r \cup i : \mathcal{X}_\sigma & \text{if } h = r, \sigma, \end{cases}$$

where $i : \mathcal{X}_\sigma = \{i : X \mid X \in \mathcal{X}_\sigma\}$ and $I(i : X) = I(X)$. We call $\gamma \in I(\mathcal{X}_h)$ a *context* for h . Each of the above sequences of sets of states is of the form $\Phi_{\gamma_\varepsilon}(\varepsilon), \dots, \Phi_{\gamma_h}(h)$, where $\Phi_\gamma(h)$ is inductively defined by:

$$\Phi_\gamma(h) = \begin{cases} \Phi_{\gamma_0} & \text{if } h = \varepsilon \\ \Phi_{\gamma_\sigma}(\Phi_{\gamma_r}(r), \sigma) & \text{if } h = r, \sigma, \end{cases}$$

where γ_0 maps every $X \in \mathcal{X}_0$ to $\gamma(0 : X)$, γ_σ maps every $X \in \mathcal{X}_\sigma$ to $\gamma(i : X)$, and $\gamma_r = \gamma|_{\mathcal{X}_r}$. Note that $\Phi_\gamma(h)$ is defined iff (i) $\Phi_{\gamma_r}(r)$ is defined, and (ii) $\pi_\sigma(\Phi_{\gamma_r}(r))$.

The following notion of a compact belief state comprises a set of contexts from $I(\mathcal{X}_h)$, a probability for each such context γ , and one for the compact belief state as a whole. Formally, the *compact belief state* $b_h^c = (p_h^c, \Gamma_h^c, Pr_h^c)$ for a history h consists of a real number $p_h^c \in [0, 1]$, a set of contexts $\Gamma_h^c \subseteq I(\mathcal{X}_h)$, and a probability function Pr_h^c on Γ_h^c . It is inductively defined as follows:

- If $h = \varepsilon$, then $p_h^c = 1$, $\Gamma_h^c = I(0 : \mathcal{X}_0)$, and $Pr_h^c(\gamma) = Pr_0(\gamma_0)$ for all $\gamma \in \Gamma_h^c$.

- If $h=r, \sigma$, then b_h^c is defined iff (i) $b_r^c=(p_r^c, \Gamma_r^c, Pr_r^c)$ is defined, and (ii) $\pi_\sigma(\Phi_\gamma(r))$ for some $\gamma \in \Gamma_r^c$. If b_h^c is defined, then it is given as follows (where i is the action length of h):

$$\begin{aligned} p_h^c &= p_r^c \cdot \sum_{\gamma \in \Gamma_r^c, \pi_\sigma(\Phi_\gamma(r))} Pr_r^c(\gamma) \\ \Gamma_h^c &= \{\gamma \cup \gamma' \mid \gamma \in \Gamma_r^c, \pi_\sigma(\Phi_\gamma(r)), \gamma' \in I(i: \mathcal{X}_\sigma)\} \\ Pr_h^c(\gamma) &= \frac{p_r^c}{p_h^c} \cdot Pr_r^c(\gamma_r) \cdot Pr_\sigma(\gamma_\sigma) \quad (\text{for all } \gamma \in \Gamma_h^c). \end{aligned}$$

Example 3.3 In the running example, Γ_ε^c comprises the contexts $(a, a), (a, b), (a, c), (b, a), (b, c), (b, b), (c, a), (c, b), (c, c)$, where every (l, l') stands for the mapping $(0: c_{at(o_1)}, 0: c_{at(o_2)}) \mapsto (l, l')$. They are under Pr_ε^c mapped to 0.03, 0.08, 0.01, 0.24, 0.48, 0.08, 0.03, 0.06, 0.01. \square

Every compact belief state $b_h^c = (p_h^c, \Gamma_h^c, Pr_h^c)$ represents the following belief state $\rho(b_h^c) = (p, \mathcal{S}, Pr)$:

$$\begin{aligned} p &= p_h^c \\ \mathcal{S} &= \{\Phi_\gamma(h) \mid \gamma \in \Gamma_h^c\} \\ Pr(S) &= \sum_{\gamma \in \Gamma_h^c, S = \Phi_\gamma(h)} Pr_h^c(\gamma) \quad (\text{for all } S \in \mathcal{S}). \end{aligned}$$

The following theorem shows that, as desired, the compact belief state b_h^c correctly implements the belief state b_h . This can be exploited for implementation (see Section 6).

Theorem 3.4 (Representation Theorem) *Let h be a history. Then, (i) the compact belief state b_h^c is defined iff the belief state b_h is defined, and (ii) $\rho(b_h^c) = b_h$.*

4 Reasoning and Planning in PC+

In this section, we show at the examples of prediction, postdiction, and planning [19, 12, 7] how important problems in probabilistic reasoning about actions can be formulated in our framework of PC+. We define them in terms of probabilities of histories in PC+. Recall that the probability of a history h , denoted $Pr(h)$, is defined as p_h , which coincides with p_h^c , by Theorem 3.4, where $b_h = (p_h, \mathcal{S}_h, Pr_h)$ is the belief state and $b_h^c = (p_h^c, \Gamma_h^c, Pr_h^c)$ is the compact belief state for h .

4.1 Prediction

We consider the following *probabilistic prediction* (or also *probabilistic temporal projection*) problem: Compute the probability that a sequence σ of actions and observations is *certainly* possible, given that another sequence σ' of actions and observations has occurred. Here, the probability that σ is certainly possible after σ' is the tight lower bound for the probability that σ is possible after σ' . We now define this probability in terms of probabilities of histories. Note that further semantically meaningful probabilities can be defined in a similar way, for example, the probability that σ *may be* possible after σ' , which is the tight upper bound for the probability that σ is possible after σ' .

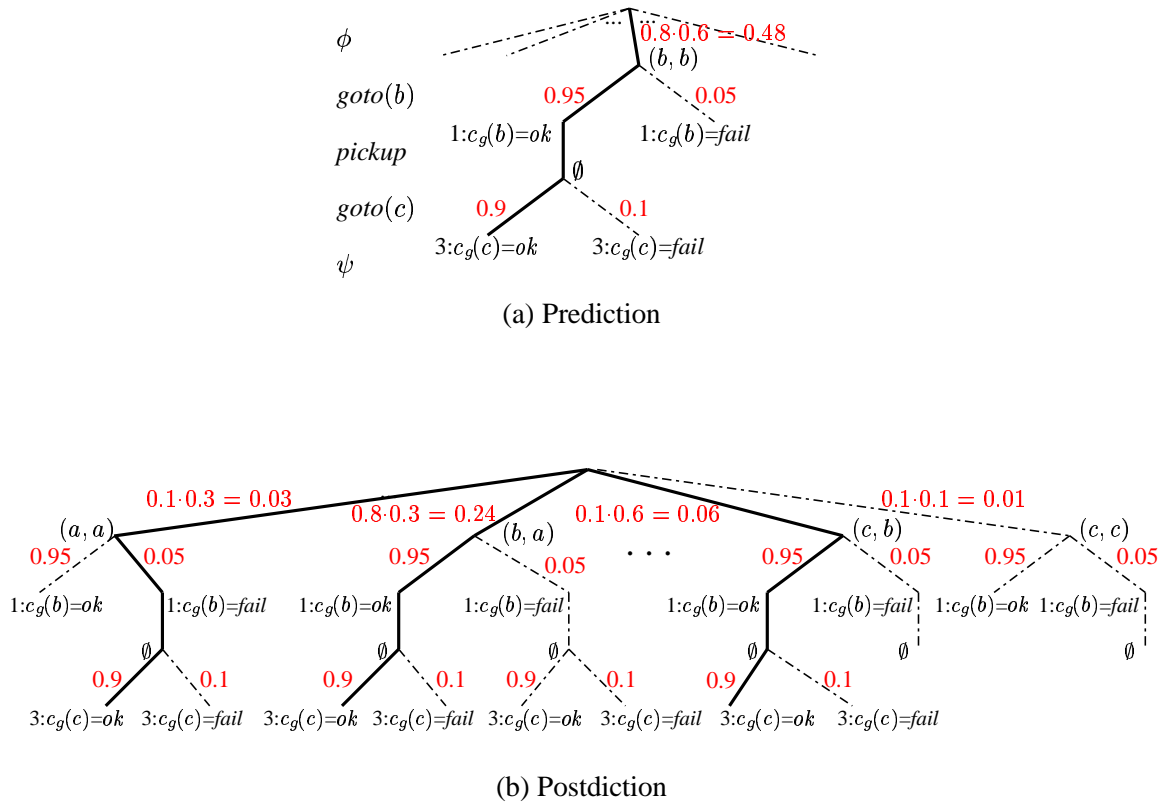


Figure 2: Evolving Contexts¹

Definition 4.1 Let $\sigma = \sigma_1, \dots, \sigma_m$ and $\sigma' = \sigma'_1, \dots, \sigma'_n$ be sequences of actions and observations. The *probability that σ is certainly possible*, denoted $Pred(\varepsilon \triangleright \sigma)$, is defined as $Pr(\Box \sigma_1, \dots, \Box \sigma_m)$. The *probability that σ is certainly possible after σ'* , denoted $Pred(\sigma' \triangleright \sigma)$, is defined as $Pr(\Diamond \sigma'_1, \dots, \Diamond \sigma'_n, \Box \sigma_1, \dots, \Box \sigma_m) / Pr(\Diamond \sigma'_1, \dots, \Diamond \sigma'_n)$.

Proposition 4.2 Let $\alpha = \alpha_1, \alpha_2, \dots, \alpha_n$ (resp., $\omega_1, \omega_2, \dots, \omega_n$) be a sequence of actions (resp., observations), and let ϕ and ψ be observations. Then,

- $Pred(\varepsilon \triangleright \phi, \alpha, \psi)$ is the probability that ϕ certainly holds initially, that α can certainly be executed, and that ψ certainly holds after that.
- $Pred(\phi \triangleright \alpha, \psi)$ is the probability that α can certainly be executed, and that then ψ certainly holds, given that ϕ is observed initially.
- $Pred(\phi, \alpha \triangleright \psi)$ is the probability that ψ certainly holds, given ϕ is observed initially, and α has been executed.

¹Pairs (x, y) are short for $(0 : c_{at(o_1)}, 0 : c_{at(o_2)}) = (x, y)$.

- $Pred(\varepsilon \triangleright \phi, \alpha_1, \omega_1, \alpha_2, \omega_2, \dots, \alpha_n, \omega_n)$ is the probability that ϕ certainly holds initially, that $\alpha = \alpha_1, \alpha_2, \dots, \alpha_n$ can certainly be executed, and that then $\omega_1, \omega_2, \dots, \omega_n$, respectively, certainly hold.

Example 4.3 Let D_0 and D be given as in Example 2.6. Suppose that the robot r is initially at location a and holds no object. Moreover, assume that the two objects o_1 and o_2 are both at location b . Then, the probability that the robot r can certainly move to b , can certainly pickup an object, and can certainly move to c , and that then at least one object is certainly at location c is given by 0.855: Let $\phi = at(r) = a \wedge at(o_1) = b \wedge at(o_2) = b \wedge holds = nil$, $\alpha = goto(b), pickup, goto(c)$, and $\psi = at(o_1) = c \vee at(o_2) = c$. We then have $Pred(\phi \triangleright \alpha, \psi) = 0.855$, which is obtained as follows (cf. Fig. 2.a). Observation ϕ leads to a context $(0 : c_{at(o_1)}, 0 : c_{at(o_2)}) = (b, b)$ of probability $0.8 \cdot 0.6$. Action $goto(b)$ then yields two contexts, $1 : c_g(b) = ok$ and $1 : c_g(b) = fail$ of probabilities 0.95 and 0.05, respectively. Action $pickup$ is only executable in $1 : c_g(b) = ok$, which leads to the empty context \emptyset of probability 1. Finally, action $goto(c)$ yields two contexts $3 : c_g(c) = ok$ and $3 : c_g(c) = fail$ of probabilities 0.9 and 0.1, respectively; in the former, ψ is true in all states, but not in the latter. In summary, $Pred(\phi \triangleright \alpha, \psi) = Pr(\diamond\phi, \square\alpha, \square\psi) / Pr(\diamond\phi) = (0.8 \cdot 0.6 \cdot 0.95 \cdot 0.9) / (0.8 \cdot 0.6) = 0.855$.

The probability that at least one object is certainly at location c after r moved to b , picked up an object, and then moved to c is given by $Pred(\phi, \alpha \triangleright \psi) = 0.9$.

The probability that r can certainly move to b , can certainly pickup an object, and can certainly move to c , and that then object o_1 is certainly at location c is given by $Pred(\phi \triangleright \alpha, \psi') = 0$, where $\psi' = at(o_1) = c$, as in no emerging context, ψ' is true in all states. Indeed, in the worst case, r might pick up o_2 and carry it to c . Thus, the probability that o_1 is at c after ϕ, α may be 0. In contrast, if we have a probability distribution for $pickup$, then $Pred(\phi \triangleright \alpha, \psi') > 0$, if r picks up o_1 with a positive probability.

Finally, suppose that object o_2 is initially at either a or b , rather than at b . Hence, consider now $\phi' = at(r) = a \wedge at(o_1) = b \wedge (at(o_2) = a \vee at(o_2) = b) \wedge holds = nil$, which is weaker than ϕ . We then obtain $Pred(\phi' \triangleright \alpha, \psi) = Pred(\phi \triangleright \alpha, \psi)$ and $Pred(\phi' \triangleright \alpha, \psi') > Pred(\phi \triangleright \alpha, \psi')$. Here, we have a positive probability, since there exists a combined context of positive probability that satisfies ϕ' , where object o_2 is absent from b , which makes $pickup$ deterministic, and thus o_1 is certainly carried to c . \square

4.2 Postdiction

Informally, the *probabilistic postdiction* (or also *probabilistic explanation*) problem that we consider here can be formulated as follows: Compute the probability that observations were certainly holding along a sequence of actions and observations ν that actually happened.

Definition 4.4 Let ν_2 be a sequence of actions and observations, and let ν_1 result from ν_2 by removing some observations. Then, the probability that ν_2 certainly occurred if ν_1 has occurred, denoted $Post(\nu_2 | \nu_1)$, is defined as $Pr(\nu_2') / Pr(\nu_1')$, where ν_2' and ν_1' result from ν_2 and ν_1 , respectively, by labeling removed observations in ν_2 with \square and all other actions and observations with \diamond .

Proposition 4.5 Let $\alpha = \alpha_1, \alpha_2, \dots, \alpha_n$ (resp., $\omega_1, \omega_2, \dots, \omega_n$) be a sequence of actions (resp., observations), $\nu = \alpha_1, \omega_1, \alpha_2, \omega_2, \dots, \alpha_n, \omega_n$, and ϕ, ψ be observations. Then,

- $Post(\phi, \alpha, \psi | \alpha, \psi)$ is the probability that ϕ certainly held initially, given α was executed and ψ was observed.
- $Post(\phi, \nu | \nu)$ is the probability that ϕ was certainly holding initially, given $\alpha_1, \alpha_2, \dots, \alpha_n$ was executed, and $\omega_1, \omega_2, \dots, \omega_n$, respectively, was observed after that.

- $Post(\phi, \nu | \alpha)$ is the probability that ϕ was certainly holding initially, and $\omega_1, \omega_2, \dots, \omega_n$ was certainly holding after $\alpha_1, \alpha_2, \dots, \alpha_n$, respectively, given α was executed.

Example 4.6 Let D_0 and D be given as in Example 2.6. Suppose that the robot r moved to b , picked up an object, and then moved to c , and that the object o_1 was observed at c after that. Then, the probability that the object o_1 was certainly at b in the initial situation is given by 0.923: Consider $\alpha = goto(b), pickup, goto(c)$, $\psi = at(o_1)=c$, and $\phi = at(o_1)=b$. We then have $Post(\phi, \alpha, \psi | \alpha, \psi) = 0.923$, which is obtained as follows: Each of the 9 initial contexts, given by the value pairs for $(0:c_{at(o_1)}, 0:c_{at(o_2)})$, except (c, c) admits execution of α resulting in a context in which ψ is observable. In detail, (b, a) , (b, b) , (b, c) , and (c, b) can be extended by $1 : c_g(b) = ok, \emptyset$, and $3 : c_g(c) = ok$, and (a, a) , (a, b) , (a, c) , and (c, a) by $1 : c_g(b) = fail, \emptyset, 3 : c_g(c) = ok$ (cf. Fig. 2.b). Let $\nu_1 = \alpha, \psi$ and $\nu_2 = \phi, \alpha, \psi$. Then, $Pr(\nu'_1) = (0.24 + 0.48 + 0.08 + 0.06) \cdot 0.95 \cdot 0.9 + (0.03 + 0.06 + 0.01 + 0.03) \cdot 0.05 \cdot 0.9 = 0.741$ and $Pr(\nu'_2) = (0.24 + 0.48 + 0.08) \cdot 0.95 \cdot 0.9 = 0.684$, as ϕ only holds in (b, a) , (b, b) , and (b, c) . Hence, $Post(\phi, \alpha, \psi | \alpha, \psi) = Pr(\nu'_2) / Pr(\nu'_1) = 0.923$. Note that if $at(r) \neq b$ would be observed after $goto(b)$ in α , then we could conclude that initially ϕ has the probability 0, which is intuitive. \square

4.3 Planning

We now formulate the notions of a (sequential) plan and of its goodness for reaching a goal observation given that a sequence of actions and observations has occurred.

Definition 4.7 Let ν be a sequence of actions and observations and ψ an observation. The sequence of actions $\alpha = \alpha_1, \dots, \alpha_n$ is a plan of goodness g for ψ after ν has occurred, denoted $Plan(\nu; \alpha; \psi) = g$, iff $Pred(\nu \triangleright \alpha, \psi) = g$.

In the general planning problem in our framework, we are then given a sequence of actions and observations ν that has occurred, a goal observation ψ , and a threshold θ , and we want to compute a plan α such that $Plan(\nu; \alpha; \psi) \geq \theta$.

Example 4.8 Let D_0 and D be as in Example 2.6. Let $\phi = at(r)=a \wedge at(o_1)=b \wedge at(o_2)=b$ and $\psi = at(o_1)=c \vee at(o_2)=c$. Then, $\alpha = goto(b), pickup, goto(c)$ is a plan of goodness 0.885 for ψ given that ϕ holds initially.

On the other hand, $Plan(\phi; \alpha; \psi') = 0$ for $\psi' = at(o_1)=c$, since r might (unpredictably) carry o_2 to c instead of o_1 . However, $Plan(\phi, \alpha', \psi') = 0.885^2 = 0.731$ for $\alpha' = \alpha, drop, goto(b), pickup, goto(c)$. Note that α' is optimal, since moving twice to c and to b is necessary in general.

If $pickup$ would be probabilistic and, e.g., obey the uniform distribution, then $Plan(\phi; \alpha; \psi') > 0$ would hold. Indeed, there would be a context c after executing α where ψ holds in all its associated states. A wrong $pickup$ decreases the success probability, which is, however, still non-zero. \square

5 Conditional Planning in PC+

In this section, we show how conditional planning can be expressed in PC+. We define the concept of a conditional plan and its goodness for reaching a goal observation. Intuitively, conditional plans generalize plans by interleaved observations, which may be used to select the proper actions for all contingencies.

Definition 5.1 A conditional plan CP is either the empty conditional plan, denoted λ , or of the form $\alpha(\omega_1 : CP_1 | \dots | \omega_n : CP_n)$, where (i) α is an action, (ii) $\Omega = \{\omega_1, \dots, \omega_n\}$, $n \geq 1$, is a set of exhaustive and pairwise exclusive observations, and (iii) CP_i , $1 \leq i \leq n$, is a conditional plan. A path for CP is either (i) the empty path, denoted ε , or (ii) of the form α or $\alpha\omega_i$, if $CP = \alpha(\omega_1 : CP_1 | \dots | \omega_n : CP_n)$, or (iii) of the form $\alpha\omega_i p_i$, if $CP = \alpha(\omega_1 : CP_1 | \dots | \omega_n : CP_n)$ and p_i is a non-empty path for CP_i .

Observe that every path p for CP addresses a subexpression of CP , denoted $CP.p$, which is defined by:

- If $p = \varepsilon$, then $CP.p = CP$.
- If $p = \alpha$ and $CP = \alpha(\omega_1 : CP_1 | \dots | \omega_n : CP_n)$, then $CP.p = (\omega_1 : CP_1 | \dots | \omega_n : CP_n)$.
- If $p = \alpha\omega_i$ and $CP = \alpha(\omega_1 : CP_1 | \dots | \omega_n : CP_n)$, then $CP.p = CP_i$.
- If $p = \alpha\omega_i p_i$ and $CP = \alpha(\omega_1 : CP_1 | \dots | \omega_n : CP_n)$, then $CP.p = CP_i.p_i$.

We next define the concept of a belief tree for a conditional plan under an initial observation. Intuitively, a belief tree is a directed tree (that is, a directed acyclic graph in which every node has exactly one parent, except for one node, called *root*, which has no parents; nodes without children are *leaves*) of possible sets of states. Its branchings represent different possible effects of actions or different possible observations. Every arc has a real number that represents the probability of its transition.

Definition 5.2 The belief tree $T = (G, Pr)$ for a conditional plan CP under an initial observation ϕ , denoted $T_{\phi CP}$, consists of a directed tree $G = (V, E)$, where the nodes are pairs (p, S) consisting of a path p for CP and a set of states S , and a mapping $Pr: E \rightarrow [0, 1]$, which are constructed by the following steps (1)–(3) (where $p \circ \beta = \beta$, if $p = \varepsilon$, and $p \circ \beta = p\beta$, otherwise):

- (1) Initially, V consists of all the arcs $(\varepsilon, \emptyset) \rightarrow (\varepsilon, S)$ such that $S = \{s \in \Phi_\gamma \mid s \models \phi\} \neq \emptyset$ for some $\gamma \in I(\mathcal{X}_0)$, and Pr maps every such arc to $\sum_{\gamma \in I(\mathcal{X}_0), S = \Phi_\gamma} Pr_0(\gamma) / \sum_{\gamma \in I(\mathcal{X}_0), \exists s \in \Phi_\gamma: s \models \phi} Pr_0(\gamma)$.
- (2) Consider the tree $T = (G, Pr)$ with $G = (V, E)$ constructed thus far. For every leaf (p, S) such that $CP.p = \alpha(\omega_1 : CP_1 | \dots | \omega_n : CP_n)$ and $\forall s \in S: s \cup \alpha \models \pi_\alpha$, we enlarge T as follows:
 - (2.1) We add to T all the arcs $(p, S) \rightarrow (p \circ \alpha, S')$ such that $S' = \Phi_\gamma(S, \alpha)$ for some $\gamma \in I(\mathcal{X}_\alpha)$, and each such arc is under Pr associated with $\sum_{\gamma \in I(\mathcal{X}_\alpha), S' = \Phi_\gamma(S, \alpha)} Pr_\alpha(\gamma)$.
 - (2.2) We then add to T all the arcs $(p \circ \alpha, S') \rightarrow (p \circ \alpha\omega_i, S'')$ with $S'' = \{s \in S' \mid s \models \omega_i\} \neq \emptyset$ along with the value 1 under Pr for each such arc.
- (3) We repeat (2) until T does not contain any leaves (p, S) such that $CP.p = \alpha(\omega_1 : CP_1 | \dots | \omega_n : CP_n)$ and $\forall s \in S: s \cup \alpha \models \pi_\alpha$.

We finally define the goodness of a conditional plan CP for obtaining a goal ψ under an initial observation ϕ . It is computed using the belief tree $T_{\phi CP} = ((V, E), Pr)$. Roughly, the success (resp., failure) leaves of $T_{\phi CP}$ are associated with the goodness value 1 (resp., 0). We then propagate such goodness values to every node of $T_{\phi CP}$, using the goodness values of the children and the numerical values associated with every arc. The goodness of CP is then defined as the goodness value of the root of $T_{\phi CP}$.

Definition 5.3 The goodness of a conditional plan CP for obtaining a goal ψ given an initial observation ϕ , denoted $CPlan(\phi; CP; \psi)$, is defined as the value $g((\varepsilon, \emptyset))$, where $g: V \rightarrow [0, 1]$ is recursively defined as follows on the set V of nodes of $T_{\phi CP} = ((V, E), Pr)$:

- $g(v) = 1$ for every leaf $v = (p, S) \in V$ (called *success leaf*) such that $CP.p = \varepsilon$ and $\forall s \in S: s \models \psi$;
- $g(v) = 0$ for every leaf $v = (p, S) \in V$ (called *failure leaf*) such that either $CP.p \neq \varepsilon$ or $\exists s \in S: s \not\models \psi$;
- $g(v) = \min_{v \rightarrow v' \in E} Pr(v \rightarrow v') \cdot g(v')$ for every non-leaf node $v = (p \alpha, S) \in V$;
- $g(v) = \sum_{v \rightarrow v' \in E} Pr(v \rightarrow v') \cdot g(v')$ for every other node $v = (p, S) \in V$.

We illustrate the concept of a conditional plan and its goodness in our robot example, where α , CP_1 and α abbreviate $\alpha (\top : CP_1)$ and $\alpha (\top : \lambda)$, respectively.

Example 5.4 For reaching the goal $\psi' = at(o_1) = c$ given the initial observation $\phi = at(r) = a \wedge at(o_1) = b \wedge at(o_2) = b$, consider the conditional plan

$$CP = goto(b), pickup(holds=o_1 : goto(c) \mid holds \neq o_1 : goto(a), drop, goto(b), pickup, goto(c)).$$

It is not difficult to verify that its goodness is given by $CPlan(\phi; CP; \psi') = 0.95^3 \cdot 0.9 = 0.772$, which is higher than $Plan(\phi; \alpha'; \psi') = 0.731$ in Example 4.8.

Clearly, $0.95 \cdot 0.9$ is an upper bound for the goodness of any (conditional or normal) plan for reaching the goal ψ' given the initial observation ϕ . The one of the conditional plan $\tau_n = goto(b)$, τ'_n converges to it ($n \rightarrow \infty$), where $\tau'_1 = pickup(holds=o_1 : goto(c) \mid holds \neq o_1 : goto(a), drop, goto(b), pickup, goto(c))$ and $\tau'_n = pickup(holds=o_1 : goto(c) \mid holds \neq o_1 : drop, \tau'_{n-1})$ for $n > 1$. \square

6 Reduction to Causal Theories

Section 4 shows that several important problems in reasoning about actions can be formulated in terms of probabilities of histories. To compute such probabilities, it is sufficient to compute compact belief states. Here, the main computational task is to decide whether $\pi_\sigma(\Phi_\gamma(h))$ holds for contexts $\gamma \in I(\mathcal{X}_h)$. We now show how this can be reduced to reasoning about models of causal theories.

We first give some preparative definitions. For each non-rigid variable $X \in \mathcal{X}$ and $i \geq 0$, let $i : X$ be a new variable of the same type and domain. We denote by $i : F$ the formula obtained from a formula F by replacing each atom $X = x$ with non-rigid $X \in \mathcal{X}$ by $i : X = x$. We define

- $0 : D_0$ as the set of all $0 : X = x \Leftarrow 0 : X = x$ for each simple fluent $X \in \mathcal{X}$ and $x \in I(X)$, and all $0 : F \Leftarrow 0 : G$ for each axiom (1) in D_0 , and
- $i : D$ as all $i : F \Leftarrow i : G$ for each axiom (1) in D , and all $i : F \Leftarrow (i : G) \wedge (i-1 : H)$ for each axiom (2) in D , $i > 0$.

We next associate every history h and context $\gamma \in I(\mathcal{X}_h)$ with a causal theory D_h^γ as follows.

Definition 6.1 For histories h of action length m , the associated causal theory D_h is inductively defined by:

$$D_h = \begin{cases} 0 : D_0 & \text{if } h = \varepsilon \\ D_r \cup m : D \cup \{m-1 : \hat{\alpha} \Leftarrow \top\} & \text{if } h = r, \circ \alpha \\ D_r \cup \{m : \omega \Leftarrow \top\} & \text{if } h = r, \circ \omega, \end{cases}$$

where α is an action, $\hat{\alpha}$ is a formula encoding α , and ω is an observation. For each $\gamma \in I(\mathcal{X}_h)$, let D_h^γ be obtained from D_h by replacing each $F \Leftarrow G$ by $\gamma(F \Leftarrow G)$.

The following theorem then finally describes how deciding whether $\pi_\sigma(\Phi_\gamma(h))$ holds can be reduced to reasoning about the models of the causal theory D_h^γ .

Theorem 6.2 *Let h be a history of action length m . Let $\gamma \in I(\mathcal{X}_h)$ and τ be an action α (resp., observation ω). Suppose that $\Phi_\gamma(h)$ is defined. Then, $\pi_{\diamond\tau}(\Phi_\gamma(h))$ iff $\exists s \in \text{Mod}(D_h^\gamma) : s \models m : \pi_\alpha$ (resp., $s \models m : \omega$), and $\pi_{\square\tau}(\Phi_\gamma(h))$ iff $\forall s \in \text{Mod}(D_h^\gamma) : s \models m : \pi_\alpha$ (resp., $s \models m : \omega$), where $\text{Mod}(D_h^\gamma)$ denotes the set of all models of D_h^γ .*

Using this result, we can inductively test if $\Phi_\gamma(h)$ is defined along the prefixes of h . As D_h^γ is polynomially constructible, we thus can compute $Pr(h)$ with an oracle for causal theory consistency (which is Σ_2^p -complete [7]) using a simple (in worst case exponential) backtracking procedure. Moreover, deciding $Pr(h) > 0$ is feasible in $\text{NP}^{\Sigma_2^p}$, i.e., in nondeterministic polynomial time with a Σ_2^p oracle.

7 Summary and Outlook

We have presented the probabilistic action language $\text{PC}+$, which generalizes $\mathcal{C}+$ by probabilistic information into an expressive framework for dealing with qualitative as well as quantitative uncertainty about actions. Its formal semantics is defined in terms of probabilistic transitions between sets of states. We have then shown that, using the concepts of a history and its belief state, several important problems in reasoning about actions can be concisely expressed in it, which may be solved employing reductions to causal theories and solvers for them.

An interesting topic of future research is to provide more efficient algorithms and a detailed complexity analysis for probabilistic reasoning about actions in $\text{PC}+$. Other interesting topics are to add costs to planning and conditional planning and to define in our framework further semantic notions like counterfactuals, interventions, actual cause, and causal explanations, taking inspiration by similar concepts in the structural-model approach to causality [15, 9].

A Proofs

Proof of Proposition 3.2. Let $h = r, \sigma$ be a history, where $\sigma = \circ \omega$ is a labeled observation. It is sufficient to show that if $b_h = (p_h, \mathcal{S}_h, Pr_h)$ is defined, then (19) holds. Thus, suppose that b_h is defined. Then,

$$\begin{aligned} p_h &= p_r \cdot \sum_{S \in \mathcal{S}_r, \pi_\sigma(S)} Pr_r(S); \\ \mathcal{S}_h &= \{\Phi_\gamma(S, \sigma) \mid \gamma \in I(\mathcal{X}_\sigma), S \in \mathcal{S}_r, \pi_\sigma(S)\} \\ &= \{\{s \in \mathcal{S} \mid s \models \omega\} \mid S \in \mathcal{S}_r, \pi_\sigma(S)\}; \\ Pr_h(S') &= \frac{p_r}{p_h} \cdot \sum_{\substack{S \in \mathcal{S}_r, \\ \pi_\sigma(S)}} Pr_\sigma(S'|S) \cdot Pr_r(S) \quad (\text{for all } S' \in \mathcal{S}_h) \\ &= \frac{p_r}{p_h} \cdot \sum_{S' = \{s \in \mathcal{S} \mid s \models \omega\}} Pr_r(S) \quad (\text{for all } S' \in \mathcal{S}_h). \end{aligned}$$

This shows that (19) holds. \square

Proof of Theorem 3.4. Observe first that $\rho(b_h^c)$ is defined, if b_h^c is defined. It is thus sufficient to show that (i) b_h^c is defined iff b_h is defined, and (ii) if b_h^c is defined, then $\rho(b_h^c) = b_h$. We show this by induction along h .

Basis: Let $h = \varepsilon$. Then, both b_h^c and b_h are defined, and thus b_h^c is defined iff b_h is defined. That is, (i) holds. Let $b_h^c = (p_h^c, \Gamma_h^c, Pr_h^c)$, $b_h = (p_h, \mathcal{S}_h, Pr_h)$, and $\rho(b_h^c) = (p, \mathcal{S}, Pr)$. Trivially, $p = p_h^c = 1 = p_h$. Observe then that $\mathcal{S} = \{\Phi_\gamma(h) \mid \gamma \in \Gamma_h^c\} = \{\Phi_{\gamma_0} \mid \gamma \in \Gamma_h^c\} = \{\Phi_\gamma \mid \gamma \in I(\mathcal{X}_0)\} = \mathcal{S}_h$. Finally, for every $S \in \mathcal{S} = \mathcal{S}_h$:

$$\begin{aligned} Pr(S) &= \sum_{\gamma \in \Gamma_h^c, S = \Phi_\gamma(h)} Pr_h^c(\gamma) \\ &= \sum_{\gamma \in I(\mathcal{X}_0), S = \Phi_\gamma} Pr_0(\gamma) \\ &= Pr_h(S). \end{aligned}$$

In summary, this shows that $\rho(b_h^c) = b_h$. That is, (ii) holds.

Induction: Let r be a history, and suppose that (i) b_r^c is defined iff b_r is defined, and (ii) if b_r^c is defined, then $\rho(b_r^c) = b_r$. We now prove that (i) b_h^c is defined iff b_h is defined, and (ii) if b_h^c is defined, then $\rho(b_h^c) = b_h$, where $h = r, \sigma$ with $\sigma = \circ \tau$. We first prove that (i) holds. By the induction hypothesis,

$$\begin{aligned} b_h^c \text{ is defined} &\text{ iff (a) } b_r^c \text{ is defined, and (b) } \pi_\sigma(\Phi_\gamma(r)) \text{ for some } \gamma \in \Gamma_r^c \\ &\text{ iff (a) } b_r \text{ is defined, and (b) } \pi_\sigma(S) \text{ for some } S \in \mathcal{S}_r \\ &\text{ iff } b_h \text{ is defined.} \end{aligned}$$

This shows that (i) holds. For the proof of (ii), assume now that b_h^c is defined. Let $b_r^c = (p_r^c, \Gamma_r^c, Pr_r^c)$ and $b_r = (p_r, \mathcal{S}_r, Pr_r)$. Let $b_h^c = (p_h^c, \Gamma_h^c, Pr_h^c)$, $b_h = (p_h, \mathcal{S}_h, Pr_h)$, and $\rho(b_h^c) = (p, \mathcal{S}, Pr)$, where $h = r, \sigma$ with $\sigma = \circ \tau$. We then obtain the following by the induction hypothesis:

$$\begin{aligned} \mathcal{S} &= \{\Phi_\gamma(h) \mid \gamma \in \Gamma_h^c\} \\ &= \{\Phi_{\gamma_\sigma}(\Phi_{\gamma_r}(r), \sigma) \mid \gamma \in \Gamma_h^c\} \\ &= \{\Phi_{\gamma_\sigma}(\Phi_{\gamma_r}(r), \sigma) \mid \gamma_r \in \Gamma_r^c, \pi_\sigma(\Phi_{\gamma_r}(r)), \gamma_\sigma \in I(\mathcal{X}_\sigma)\} \\ &= \{\Phi_\gamma(S, \sigma) \mid \gamma \in I(\mathcal{X}_\sigma), S \in \mathcal{S}_r, \pi_\sigma(S)\} \\ &= \mathcal{S}_h. \end{aligned}$$

Moreover, the following holds by the induction hypothesis:

$$\begin{aligned} p &= p_h^c \\ &= p_r^c \cdot \sum_{\gamma \in \Gamma_r^c, \pi_\sigma(\Phi_\gamma(r))} Pr_r^c(\gamma) \\ &= p_r \cdot \sum_{S \in \mathcal{S}_r, \pi_\sigma(S), \gamma \in \Gamma_r^c, S = \Phi_\gamma(r)} Pr_r^c(\gamma) \\ &= p_r \cdot \sum_{S \in \mathcal{S}_r, \pi_\sigma(S)} Pr_r(S) \\ &= p_h. \end{aligned}$$

Finally, for every $S \in \mathcal{S} = \mathcal{S}_h$, the following holds by the induction hypothesis:

$$\begin{aligned}
Pr(S) &= \sum_{\gamma \in \Gamma_h^c, S = \Phi_\gamma(h)} Pr_h^c(\gamma) \\
&= \frac{p_r^c}{p_h^c} \cdot \sum_{\gamma \in \Gamma_h^c, S = \Phi_\gamma(h)} Pr_r^c(\gamma_r) \cdot Pr_\sigma(\gamma_\sigma) \\
&= \frac{p_r^c}{p_h^c} \cdot \sum_{\gamma_r \in \Gamma_r^c, \pi_\sigma(\Phi_{\gamma_r}(r)), \gamma_\sigma \in I(\mathcal{X}_\sigma), S = \Phi_{\gamma_\sigma}(\Phi_{\gamma_r}(r), \sigma)} Pr_r^c(\gamma_r) \cdot Pr_\sigma(\gamma_\sigma) \\
&= \frac{p_r}{p_h} \cdot \sum_{S' \in \mathcal{S}_r, \pi_\sigma(S'), \gamma_r \in \Gamma_r^c, S' = \Phi_{\gamma_r}(r)} Pr_r^c(\gamma_r) \cdot (\sum_{\gamma_\sigma \in I(\mathcal{X}_\sigma), S = \Phi_{\gamma_\sigma}(S', \sigma)} Pr_\sigma(\gamma_\sigma)) \\
&= \frac{p_r}{p_h} \cdot \sum_{S' \in \mathcal{S}_r, \pi_\sigma(S')} \sum_{\gamma \in \Gamma_r^c, S' = \Phi_\gamma(r)} Pr_r^c(\gamma) \cdot (\sum_{\gamma \in I(\mathcal{X}_\sigma), \Phi_\gamma(S', \sigma) = S} Pr_\sigma(\gamma)) \\
&= \frac{p_r}{p_h} \cdot \sum_{S' \in \mathcal{S}_r, \pi_\sigma(S')} Pr_r(S') \cdot (\sum_{\gamma \in I(\mathcal{X}_\sigma), \Phi_\gamma(S', \sigma) = S} Pr_\sigma(\gamma)) \\
&= \frac{p_r}{p_h} \cdot \sum_{S' \in \mathcal{S}_r, \pi_\sigma(S')} Pr_r(S') \cdot Pr_\sigma(S|S') \\
&= Pr_h(S).
\end{aligned}$$

In summary, we obtain $\rho(b_h^c) = b_h$. This shows that (ii) holds. \square

Proof of Propositions 4.2 and 4.5. The statements follow immediately from the observation that actually observed formulas (resp., preconditions of actually executed actions) must hold in at least one element of a possible set of states, while hypothetically observed formulas (resp., preconditions of hypothetical actions) must hold in all elements of a possible set of states. \square

For the proof of Theorem 6.2, we need the following preparative result.

Proposition A.1 *Let h be a history of action length m , let $\gamma \in I(\mathcal{X}_h)$, and let α (resp., ω) be an action (resp., observation). Suppose $\Phi_\gamma(h)$ is defined. Then, $Mod(D_h^\gamma)$ is given by*

$$\begin{cases}
\{0 : s \mid s \in \Phi_\gamma(h)\} & \text{if } h = \varepsilon \\
\{t \in Mod(D_r^{\gamma_r}) \mid t \models m : \omega\} & \text{if } h = r, \circ \omega \\
\{t \cup (m : s) \cup (m-1 : \alpha) \mid t \in Mod(D_r^{\gamma_r}), s \in \Phi_{\gamma_\alpha}(t_{m-1}, \alpha)\} & \text{if } h = r, \circ \alpha,
\end{cases}$$

where (i) $m : s$ maps every fluent variable $m : X$ with $X \in \mathcal{X}$ to $s(X)$, (ii) $m-1 : \alpha$ maps every action variable $m-1 : X$ with $X \in \mathcal{X}$ to $\alpha(X)$, and (iii) t_{m-1} maps every rigid (resp., fluent) variable $X \in \mathcal{X}$ to $t(X)$ (resp., $t(m-1 : X)$).

Proof. For $h = \varepsilon$, we have $D_h^\gamma = (0 : D_0)^\gamma$ and $\Phi_\gamma(h) = \Phi_{\gamma_0}$, where (i) $(0 : D_0)^\gamma$ is the set of all rules $\gamma(0 : F \Leftarrow 0 : G)$ for all axioms (1) in D_0 and all rules $0 : X = x \Leftarrow 0 : X = x$ for each simple fluent $X \in \mathcal{X}$ and $x \in I(X)$, and (ii) Φ_{γ_0} is the set of all models over rigid and fluent variables of the causal theory comprising the rules $\gamma_0(F \Leftarrow G)$ for all axioms (1) in D_0 and all rules $X = x \Leftarrow X = x$ for each simple fluent $X \in \mathcal{X}$ and $x \in I(X)$. This shows that $Mod(D_h^\gamma) = \{0 : s \mid s \in \Phi_\gamma(h)\}$.

For $h = r, \circ \omega$, we have $D_h^\gamma = D_r^{\gamma_r} \cup \{m : \omega \Leftarrow \top\}$ and thus $Mod(D_h^\gamma) = \{t \in Mod(D_r^{\gamma_r}) \mid t \models m : \omega\}$.

For $h = r, \circ \alpha$, it holds that $D_h^\gamma = D_r^{\gamma_r} \cup \{m : D\}^\gamma \cup \{m-1 : \hat{\alpha} \Leftarrow \top\}$. Suppose that $t' \in Mod(D_h^\gamma)$. Then, $t'(X) = (m-1 : \alpha)(X)$ for every action variable $X \in \mathcal{X}$. Moreover, the restriction t'_r of t' to all rigid variables $X \in \mathcal{X}$, all fluent variables $i : X$ such that $0 \leq i \leq m-1$ and $X \in \mathcal{X}$, and all action variables $i : X$ such that $0 \leq i \leq m-2$ and $X \in \mathcal{X}$ is a model of the reduct of $D_r^{\gamma_r}$ relative to t'_r . We now show that t'_r is also a unique such model. Towards a contradiction, suppose the contrary. That is, there exists another such

model t'' . Then, $t'' \cup t' | m \cup (m-1 : \alpha)$, where $t' | m$ is the restriction of t' to all fluent variables $m : X$ with $X \in \mathcal{X}$, is a model of the reduct of D_h^γ relative to t' , which contradicts $t' \in \text{Mod}(D_h^\gamma)$. This shows that $t'_r \in \text{Mod}(D_r^{\gamma_r})$. Finally, it holds that t'_m satisfies all formulas caused in (t'_{m-1}, α, t'_m) under γ_α . We now show that t'_m is also a unique such interpretation. Towards a contradiction, suppose there exists another such interpretation t'' . Then, $t'_r \cup t'' | m \cup (m-1 : \alpha)$ is a model of the reduct of D_h^γ relative to t' , which contradicts $t' \in \text{Mod}(D_h^\gamma)$. This shows that t'_m is a unique such interpretation, and thus $t'_m \in \Phi_{\gamma_\alpha}(t'_{m-1}, \alpha)$. In summary, $\text{Mod}(D_h^\gamma)$ is a subset of the set of all $t \cup (m : s) \cup (m-1 : \alpha)$ such that $t \in \text{Mod}(D_r^{\gamma_r})$ and $s \in \Phi_{\gamma_\alpha}(t_{m-1}, \alpha)$. To prove the converse, suppose that $t \in \text{Mod}(D_r^{\gamma_r})$ and $s \in \Phi_{\gamma_\alpha}(t_{m-1}, \alpha)$. Then, $t' = t \cup (m : s) \cup (m-1 : \alpha)$ is a model of the reduct of D_h^γ relative to t' . We now prove that t' is also a unique such model. Towards a contradiction, suppose there exists another such model t'' . Then, $t''(m-1 : X) = t(m-1 : X)$ for every action variable $X \in \mathcal{X}$. Moreover, $t''_r = t'_r$ as otherwise t''_r would not be a model of $D_r^{\gamma_r}$. Hence, $t'' | m \neq t' | m$, but this contradicts (t_{m-1}, α, s) being causally explained under γ_α . This shows that t' is a unique model of the reduct of D_h^γ relative to t' , and thus $t' \in \text{Mod}(D_h^\gamma)$. In summary, every $t \cup (m : s) \cup (m-1 : \alpha)$ such that $t \in \text{Mod}(D_r^{\gamma_r})$ and $s \in \Phi_{\gamma_\alpha}(t_{m-1}, \alpha)$ is also an element of $\text{Mod}(D_h^\gamma)$. \square

We are now ready to prove Theorem 6.2.

Proof of Theorem 6.2. It is sufficient to show that

$$\Phi_\gamma(h) = \{t_m \mid t \in \text{Mod}(D_h^\gamma)\}. \quad (20)$$

We give a proof by induction along h .

Basis: Immediate by Proposition A.1 for the case $h = \varepsilon$.

Induction: Assume that $h = r, \sigma$, where $\sigma = \circ \omega$ (resp., $\sigma = \circ \alpha$) and $n = \text{alen}(r) = \text{alen}(h) = m$ (resp., $n = \text{alen}(r) = \text{alen}(h) - 1 = m - 1$). Suppose that $\Phi_{\gamma_r}(r) = \{t_n \mid t \in \text{Mod}(D_r^{\gamma_r})\}$.

For $\sigma = \circ \omega$, the following holds by the induction hypothesis and by Proposition A.1:

$$\begin{aligned} \Phi_\gamma(h) &= \Phi_{\gamma_\sigma}(\Phi_{\gamma_r}(r), \sigma) \\ &= \{s \in \Phi_{\gamma_r}(r) \mid s \models \omega\} \\ &= \{t_m \mid t \in \text{Mod}(D_r^{\gamma_r}), t \models m : \omega\} \\ &= \{t_m \mid t \in \text{Mod}(D_h^\gamma)\}. \end{aligned}$$

For $\sigma = \circ \alpha$, the following holds by the induction hypothesis and by Proposition A.1:

$$\begin{aligned} \Phi_\gamma(h) &= \Phi_{\gamma_\sigma}(\Phi_{\gamma_r}(r), \sigma) \\ &= \{s \mid \exists s' \in \Phi_{\gamma_r}(r) : s \in \Phi_{\gamma_\alpha}(s', \alpha)\} \\ &= \{s \mid \exists t \in \text{Mod}(D_r^{\gamma_r}) : s \in \Phi_{\gamma_\alpha}(t_{m-1}, \alpha)\} \\ &= \{t_m \mid t \in \text{Mod}(D_h^\gamma)\}. \end{aligned}$$

In summary, this shows that (20) holds. \square

References

- [1] F. Bacchus, J. Y. Halpern, and H. J. Levesque. Reasoning about noisy sensors and effectors in the situation calculus. *Artif. Intell.*, 111(1-2):171–208, 1999.
- [2] C. Baral and M. Gelfond. Reasoning about effects of concurrent actions. *J. Logic Program.*, 31:85–117, 1997.
- [3] C. Baral, N. Tran, and L.-C. Tuan. Reasoning about actions in a probabilistic setting. In *Proceedings AAAI-2002*, pp. 507–512, 2002.
- [4] T. Eiter, W. Faber, N. Leone, G. Pfeifer, and A. Polleres. A logic programming approach to knowledge-state planning, II: The DLV^K system. *Artif. Intell.*, 144(1-2):157–211, 2003.
- [5] A. Finzi and F. Pirri. Combining probabilities, failures and safety in robot control. In *Proceedings IJCAI-2001*, pp. 1331–1336, 2001.
- [6] M. Gelfond and V. Lifschitz. Representing action and change by logic programs. *J. Logic Program.*, 17:301–322, 1993.
- [7] E. Giunchiglia, J. Lee, V. Lifschitz, N. McCain, and H. Turner. Nonmonotonic causal theories. *Artif. Intell.*, 2003. To appear.
- [8] E. Giunchiglia and V. Lifschitz. An action language based on causal explanation: Preliminary report. In *Proceedings AAAI-1998*, pp. 623–630, 1998.
- [9] J. Y. Halpern and J. Pearl. Causes and explanations: A structural-model approach – Part II: Explanations. In *Proceedings IJCAI-01*, pp. 27–34, 2001.
- [10] J. Y. Halpern and M. R. Tuttle. Knowledge, probability, and adversaries. *J. ACM*, 40(4):917–962, 1993.
- [11] L. P. Kaelbling, M. L. Littman, and A. R. Cassandra. Planning and acting in partially observable stochastic domains. *Artif. Intell.*, 101(1-2):99–134, 1998.
- [12] P. Mateus, A. Pacheco, and J. Pinto. Observations and the probabilistic situation calculus. In *Proceedings KR-2002*, pp. 327–338, 2002.
- [13] N. McCain and H. Turner. Causal theories of action and change. In *Proceedings AAAI-1997*, pp. 460–465, 1997.
- [14] N. McCain and H. Turner. Satisfiability planning with causal theories. In *Proceedings KR-1998*, pp. 212–223, 1998. For the Causal Calculator see <http://www.cs.utexas.edu/users/tag/cc/>.
- [15] J. Pearl. *Causality: Models, Reasoning, and Inference*. Cambridge University Press, 2000.
- [16] D. Poole. The independent choice logic for modelling multiple agents under uncertainty. *Artif. Intell.*, 94(1-2):7–56, 1997.
- [17] D. Poole. Decision theory, the situation calculus and conditional plans. *Electronic Transactions on Artificial Intelligence*, 2(1-2):105–158, 1998.
- [18] D. Poole. Logic, knowledge representation, and Bayesian decision theory. In *Proceedings CL-2000*, pp. 70–86, 2000.
- [19] R. Reiter. *Knowledge in Action: Logical Foundations for Specifying and Implementing Dynamical Systems*. MIT Press, 2001.
- [20] H. Turner. Representing actions in logic programs and default theories: A situation calculus approach. *J. Logic Program.*, 31:245–298, 1997.