Dealing with Incomplete Transition Information in a Stochastic Action Logic

Gavin Rens[1] Thomas Meyer[1]
Gerhard Lakemeyer[2]

[1]Centre for Artificial Intelligence Research, University of KwaZulu-Natal, and CSIR Meraka, South Africa
[2]RWTH Aachen University, Germany

April 22, 2014

Abstract

We investigate the requirements for specifying the behaviors of actions in a stochastic domain. That is, we propose how to write sentences in a logical language to capture a model of stochastic transitions due to the execution of actions of some agent. We propose a definition for 'proper' and 'complete' probabilistic transition model specifications and suggest which assumptions can and perhaps should be made about such specifications to make them more parsimonious. Making *a priori* or default assumptions about the nature of transitions is useful when a given transition model is not fully specified. Two default assumption approaches will be considered.

1 Introduction

Many environments can be modeled as probabilistic transition systems. For instance, a robot which is uncertain about the outcomes of its actions could rely on such a model. Or to simulate some biological process may require a model of how likely it is that a particular state of the process will arise, given some (cellular/molecular/chemical) event occurs in another process state. Usually, a full specification of transition probabilities is required so that the likelihood of the system changing from one current state s_c to a resulting state s_r can be deduced, for *all* system states.

There are naïve ways of specifying a system's dynamics and there are more sophisticated ways which attempt to make the task of specification easier and the specifications more compact, by making use of regularities and common sense. In this report, intuitive lines of reasoning are followed, relying on two kinds of default assumptions when transition information is deficient. Transition information may be unobtainable or difficult to deduce, or the knowledge engineer may know that the default assumption is correct for a given domain and thus needs not (re)state the implicit information.

Imagine a robot that is in need of an oil refill. There is an open can of oil on the floor within reach of its gripper. If there is nothing else in the robot's gripper, it can grab the can (or miss it, or knock it over) and it can drink the oil by lifting the can to its 'mouth' and pouring the contents in (or miss its mouth and spill). The robot may also want to confirm whether there is anything left in the oil-can by weighing its contents with its 'weight' sensor. And once holding the can, the robot may wish to replace it on the floor.

The domain is (partially) formalized as follows. The robot has the set of (intended) actions $\mathcal{A} = \{\mathtt{grab}, \mathtt{drink}, \mathtt{weigh}, \mathtt{replace}\}$ with expected intuitive meanings. The robot experiences its environment through three Boolean features: $\mathcal{P} = \{\mathtt{full}, \mathtt{drank}, \mathtt{holding}\}$ meaning that the oil-can is full, that the robot has drunk the oil and that it is currently holding something in its gripper. Given a formalization BK of our scenario, the robot may have the following query: If the oil-can is empty and I'm not holding it, is there a 0.9 probability that I'll be holding it (and it is still empty) after grabbing it, and a 0.1 probability that I'll have missed it (and it is still empty)? That is, does $(\neg\mathtt{full} \land \neg\mathtt{holding}) \rightarrow ([\mathtt{grab}]_{0.9}(\neg\mathtt{full} \land \mathtt{holding}) \land [\mathtt{grab}]_{0.1}(\neg\mathtt{full} \land \neg\mathtt{holding}))$ follow from BK?

Next, we define a general logic for reasoning about agents with stochastic actions. In Section 3, we introduce the basic approach to specifying a domain. Section 4 investigates how invariance of features of the world under certain conditions can be captured. The domain specification approach is then extended employing the new insights. Section 5 tackles the issue of how to complete underspecified specifications. In Section 6, we prove that our two approaches for completing specifications lead to 'full' specifications. We conclude with some remarks.

2 A Stochastic Action Logic

We use the Specification Logic of Actions with Probability (SLAP) Rens et al. [2013b] as the vehicle to convey our ideas.

The vocabulary of our language contains three sorts of objects of interest:

1. a finite set of *propositional variables* (alias, *fluents*) $\mathcal{P} = \{p_1, \ldots, p_n\}$,
2. a finite set of names of atomic *actions* $\mathcal{A} = \{\alpha_1, \ldots, \alpha_n\}$,
3. all *rational numbers* $\mathbb{Q}$.

$\mathbb{Q} \cap [0,1]$ will be denoted as $\mathbb{Q}_{[0,1]}$. We are going to work in a multi-modal setting, in which we have modal operators $[\alpha]_q$, one for each $\alpha \in \mathcal{A}$ and $q \in \mathbb{Q}_{[0,1]}$.

Definition 2.1 *Let $\alpha \in \mathcal{A}$, $q \in \mathbb{Q}_{[0,1]}$ and $p \in \mathcal{P}$. The language of SLAP, denoted $\mathcal{L}_{SLAP}$, is the least set of Ψ defined by the grammar:*

$$\begin{aligned}
\varphi &::= \; p \mid \top \mid \neg\varphi \mid \varphi \wedge \varphi. \\
\Phi &::= \; \varphi \mid \neg\Phi \mid \Phi \wedge \Phi \mid [\alpha]_q\varphi. \\
\Psi &::= \; \Phi \mid \Box\Phi \mid \Psi \wedge \Psi.
\end{aligned}$$

We shall also require the definition of $\mathcal{L}^{-\Box}_{SLAP}$, the least set of Φ as defined above.

In SLAP, sentences of the form $\neg\Box\Phi$ are not in the language. The reason is that the decision procedure for SLAP entailment would not notice certain contradictions which may occur due to such sentences being allowed. Note that formulae with nested modal operators are not in $\mathcal{L}_{SLAP}$. As usual, we treat $\bot, \vee, \rightarrow$ and $\leftrightarrow$ as abbreviations. $\rightarrow$ and $\leftrightarrow$ have the weakest bindings and $\neg$ the strongest; parentheses enforce or clarify the scope of operators conventionally.

Two distinguished schemata are $[\alpha]_q\varphi$ and $\neg[\alpha]_q\varphi$ and shall be referred to as *dynamic literals*. Any formula which includes a dynamic literal shall be referred to as *dynamic*. $[\alpha]_q\varphi$ is read 'The probability of reaching a world in which φ holds after executing α, is equal to q'. $[\alpha]$ abbreviates $[\alpha]_1$. $\langle\alpha\rangle\varphi$ abbreviates $\neg[\alpha]_0\varphi$ and is read 'It is possible to reach a world in which φ holds after executing α'. Note that $\langle\alpha\rangle\varphi$ does not mean $\neg[\alpha]\neg\varphi$. One reads $\Box\Phi$ as 'Φ holds in every possible world'. We require the $\Box$ operator to mark certain information (sentences) as holding in *all* possible worlds—essentially, the axioms which model the domain of interest.

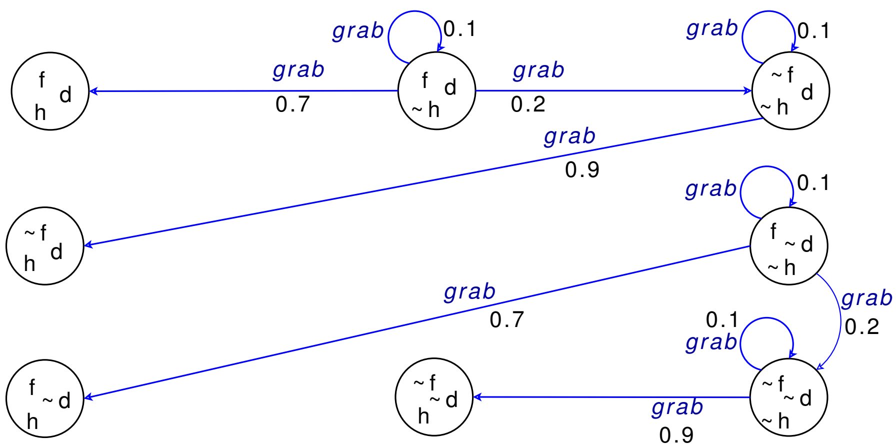

Figure 1: A transition diagram for the `grab` action. The letters f, d and h, respectively represent propositional literals `full`, `drank` and `holding`. And $\sim$ reads 'not'.

Intuitively, a world w defines a set of features (*propositions*) that the agent understands and that describes a state of affairs in the world or that describes a possible, alternative world. Let $w : \mathcal{P} \mapsto \{0, 1\}$ be a total function that assigns a truth value to each proposition. Let C be the set of all possible functions w. We call C the *conceivable worlds*.

Definition 2.2 *A SLAP structure is a tuple $\mathcal{S} = \langle W, R \rangle$ such that*

1. $W \subseteq C$ *a non-empty set of* possible worlds.
2. $R : \mathcal{A} \mapsto R_\alpha$*, where* $R_\alpha : (W \times W) \mapsto \mathbb{Q}_{[0,1]}$ *is a total function from pairs of worlds into the rationals; That is,* R *is a mapping that provides an accessibility relation* R_α *for each action* $\alpha \in \mathcal{A}$*; For every* $w^- \in W$*, it is required that either* $\sum_{(w^-,w^+,pr) \in R_\alpha} pr = 1$ *or* $\sum_{(w^-,w^+,pr) \in R_\alpha} pr = 0$.

Note that the set of possible worlds may be the whole set of conceivable worlds.

R_α defines the transition probability $pr \in \mathbb{Q}_{[0,1]}$ between worlds w^+ and w^- via action α. If $(w^-, w^+, 0) \in R_\alpha$, then w^+ is said to be *inaccessible* or *not reachable* via α performed in w^-, else if $(w^-, w^+, pr) \in R_\alpha$ for $pr \in (0, 1]$, then w^+ is said to be *accessible* or *reachable* via action α performed in w^-. If for some w^-, $\sum_{(w^-,w^+,pr) \in R_\alpha} pr = 0$, we say that α is *inexecutable* in w^-.

Figure 1 is a pictorial representation of transitions and their probabilities for the action `grab` of the oil-can scenario. The eight circles represent the eight conceivable worlds with their valuations.

Definition 2.3 (Truth Conditions) *Let $\mathcal{S}$ be a SLAP structure, with $\alpha, \alpha' \in \mathcal{A}$ and $q, pr \in \mathbb{Q}_{[0,1]}$. Let $p \in \mathcal{P}$ and let Ψ and φ be sentence in $\mathcal{L}_{SLAP}$. We say Ψ is* satisfied *at world w in structure $\mathcal{S}$ (written $\mathcal{S}, w \models \Psi$) if and only if the following holds:*

1. $\mathcal{S}, w \models \top$ for all $w \in W$;
2. $\mathcal{S}, w \models p \iff w(p) = 1$ for $w \in W$;
3. $\mathcal{S}, w \models \neg\Psi \iff \mathcal{S}, w \not\models \Psi$;
4. $\mathcal{S}, w \models \Psi \wedge \Psi' \iff \mathcal{S}, w \models \Psi$ and $\mathcal{S}, w \models \Psi'$;
5. $\mathcal{S}, w \models [\alpha]_q\varphi \iff \big(\sum_{(w,w',pr)\in R_\alpha, \mathcal{S}, w' \models \varphi} pr\big) = q$;
6. $\mathcal{S}, w \models \Box\Psi \iff$ for all $w' \in W$, $\mathcal{S}, w' \models \Psi$.

Looking at Figure 1, for instance, if the robot is in a situation where the oil-can is full, the oil has not been drunk and the can is not being held, then the probability that the oil-can is still full after grabbing the can is 0.7 + 0.1 = 0.8. Thus, in the syntax of SLAP, given a formalization BK of the scenario, $(\texttt{full} \wedge \neg\texttt{drank} \wedge \neg\texttt{holding}) \rightarrow [\texttt{grab}]_{0.8}\texttt{full}$ follows from BK.

A formula Ψ is *valid* in a SLAP structure (denoted $\mathcal{S} \models \Psi$) if $\mathcal{S}, w \models \Psi$ for every $w \in W$. Ψ is *SLAP-valid* (denoted $\models \Psi$) if Ψ is true in every structure $\mathcal{S}$. If $\models \theta \leftrightarrow \psi$, we say θ and ψ are *semantically equivalent* (abbreviated $\theta \equiv \psi$).

Ψ is *satisfiable* if $\mathcal{S}, w \models \Psi$ for some $\mathcal{S}$ and $w \in W$. A formula that is not satisfiable is *unsatisfiable* or a *contradiction*. The truth of a propositional formula depends only on the world in which it is evaluated. One may thus write $w \models \Psi$ instead of $\mathcal{S}, w \models \Psi$ when Ψ is a propositional formula.

Let $\mathcal{K} \subset \mathcal{L}_{SLAP}$ and $\Phi \in \mathcal{L}_{SLAP}$. It is said that Φ is a *local semantic consequence* of $\mathcal{K}$ (denoted $\mathcal{K} \models \Phi$) if for all structures $\mathcal{S}$ and all $w \in W$ of $\mathcal{S}$, if for all $\theta \in \mathcal{K}$, $\mathcal{S}, w \models \theta$, then $\mathcal{S}, w \models \Phi$. One says that $\mathcal{K}$ *entails* Φ whenever $\mathcal{K} \models \Phi$.

3 Specifying Domains

We provide a framework to formally specify—in the language of SLAP—the domain in which an agent or robot is expected to live. In the context of SLAP, we are interested in three things in the domain of interest: (i) The initial condition IC, (ii) static laws SL of the form $\varphi \rightarrow \varphi'$ and (iii) action rule AR of the form $\varphi \rightarrow \Phi$.

Let the union of all the axioms in SL and AR be denoted by the set BK—the agent's *background knowledge*. One is typically interested in determining whether

$$\{\Box\beta \mid \beta \in BK\} \models IC \rightarrow \Phi$$

holds, where $\Phi \in \mathcal{L}_{SLAP}$ is any sentence of interest, $IC \in \mathcal{L}^{-\Box}_{SLAP}$ and $BK \subset \mathcal{L}^{-\Box}_{SLAP}$.

In SLAP, one can express that action α has effect φ with probability q under condition ϕ as $\phi \rightarrow [\alpha]_q\varphi$. In general, an effect axiom has the form

$$\phi \rightarrow [\alpha]_{q_1}\varphi_1 \wedge [\alpha]_{q_2}\varphi_2 \wedge \ldots \wedge [\alpha]_{q_n}\varphi_n$$

for expressing the different effects of an action and their associated occurrence probabilities, under a particular condition. To set the stage, we provide a definition of a 'proper' specification of the probabilistic effects of an action.

Definition 3.1 *For some action $\alpha \in \mathcal{A}$, a set of effect axioms is a* proper effects specification *(or PES for short) if and only if it takes the form*

$$\begin{gathered}\phi_1 \rightarrow [\alpha]_{q_{11}}\varphi_{11} \wedge \cdots \wedge [\alpha]_{q_{1n}}\varphi_{1n}\\ \phi_2 \rightarrow [\alpha]_{q_{21}}\varphi_{21} \wedge \cdots \wedge [\alpha]_{q_{2n}}\varphi_{2n}\\ \vdots\\ \phi_j \rightarrow [\alpha]_{q_{j1}}\varphi_{j1} \wedge \cdots \wedge [\alpha]_{q_{jn}}\varphi_{jn},\end{gathered}$$

where (i) no $q_{ik} = 0$, (ii) the transition probabilities $q_{i1}, \ldots, q_{in}$ of any axiom i must sum to 1, (iii) for every i, for any pair of effects φ_{ik} and $\varphi_{ik'}$, $\varphi_{ik} \wedge \varphi_{ik'} \equiv \bot$ and (iv) for any pair of conditions ϕ_i and $\phi_{i'}$, $\phi_i \wedge \phi_{i'} \equiv \bot$.

We insist that no $q_{ik} = 0$, because the definition is of the specification of an action's *effects*: suppose $\phi \rightarrow \ldots \wedge [\alpha]_0\varphi \wedge \cdots$ is an axiom of our background knowledge, then due to no φ-world being reachable via α under condition ϕ, φ cannot be an effect in this case. This axiom should thus not be an *effect* axiom.

The following abbreviations for constants in our scenario are used: `grab` $:= g$, `drink` $:= di$, `replace` $:= r$, `full` $:= f$, `drank` $:= d$, `holding` $:= h$. Proper specifications of the probabilistic effects of actions `grab`, `drink` and `replace`, respectively, are

$$\begin{array}{rcl}
f \wedge d \wedge \neg h & \rightarrow & [g]_{0.7}(f \wedge d \wedge h) \wedge [g]_{0.3}(d \wedge \neg h); \\
f \wedge \neg d \wedge \neg h & \rightarrow & [g]_{0.7}(f \wedge \neg d \wedge h) \wedge [g]_{0.3}(\neg d \wedge \neg h); \\
\neg f \wedge d \wedge \neg h & \rightarrow & [g]_{0.9}(\neg f \wedge d \wedge h) \wedge [g]_{0.1}(\neg f \wedge d \wedge \neg h); \\
\neg f \wedge \neg d \wedge \neg h & \rightarrow & [g]_{0.9}(\neg f \wedge \neg d \wedge h) \wedge [g]_{0.1}(\neg f \wedge \neg d \wedge \neg h); \\
f \wedge \neg d \wedge h & \rightarrow & [di]_{0.85}(\neg f \wedge d \wedge h) \wedge [d]_{0.15}(\neg f \wedge \neg d \wedge h); \\
\neg f \wedge d \wedge h & \rightarrow & [di](\neg f \wedge d \wedge h); \\
\neg f \wedge \neg d \wedge h & \rightarrow & [di](\neg f \wedge h); \\
f \wedge d \wedge h & \rightarrow & [r](f \wedge d \wedge \neg h); \\
f \wedge \neg d \wedge h & \rightarrow & [r](f \wedge \neg d \wedge \neg h); \\
\neg f \wedge d \wedge h & \rightarrow & [r](\neg f \wedge d \wedge \neg h); \\
\neg f \wedge \neg d \wedge h & \rightarrow & [r](\neg f \wedge \neg d \wedge \neg h).
\end{array}$$

The above set of axioms will be denoted as PES_1.

When trying to capture the behavior or dynamics of an action, one typically wants to capture what objects in the environment the action affects, what objects are not affected, in what situations/conditions the action can be performed and when it can physically not be performed. Observe that action α is *executable* under condition ϕ if there exists an effect axiom with condition ϕ in a PES for α. But one cannot say—given only a PES—when α is *inexecutable* or whether the action may be executable under unmentioned conditions. If a knowledge engineer for some reason does not specify what an action α's effects are, given some condition ϕ, but s/he wants to specify that the action is executable in ϕ, then s/he can simply write $\phi \rightarrow [\alpha]_1\top$. One can also write $\phi \rightarrow \langle\alpha\rangle\top$ to express executability.

Proposition 3.1 $[\alpha]\top \equiv \langle\alpha\rangle\top$.

Proof:
Let $\mathcal{S}$ be an arbitrary SLAP structure and w a world in it.

$\mathcal{S}, w \models [\alpha]\top$

$\iff$ $\big(\sum_{(w,w',pr)\in R_\alpha, \mathcal{S},w'\models\top} pr\big) = 1$ (by the truth condition for $[\alpha]\top$)

$\iff$ there exists a $w' \in W$ s.t. $(w, w', pr) \in R_\alpha$ where $pr > 0$ (by Def. **??**)

$\iff \left(\sum_{(w,w',pr)\in R_\alpha,\mathcal{S},w'\models\top} pr\right) \neq 0$
$\iff \neg[\alpha]_0\top$
$\iff \langle\alpha\rangle\top.$ ■

To express that α is inexecutable under condition ϕ, the knowledge engineer can write $\phi \to [\alpha]_0\top$.

And a PES does not carry the information of whether the axioms are meant to cover *all* conditions. What is one to make of actions performed under conditions not mentioned?

Finally, a PES is not always completely informative with respect to effects. For instance, the first axiom of PES_1 does not tell us what the probability is of reaching a world where $f \wedge d \wedge \neg h$ is satisfied, given `grab` is performed under condition $f \wedge d \wedge \neg h$.

The rest of this report is dedicated to dealing with these deficits.

4 Invariace

A *frame axiom* Reiter [1991] captures the idea of the 'momentum' of a state. That is, things which are unaffected by an action, should remain unaffected after the completion of the action. The general problem of how to minimize or avoid specifying the multitude of frame axioms usually required is known as the frame problem McCarthy and Hayes [1969]. Bacchus, Halpern and Levesque Bacchus et al. [1999] supply one approach to deal with the frame problem in a language able to express probabilistic transitions, but read the concluding remarks at the end of this report.

We see in PES_1 that for the action r, only h is affected. So for r, the four frame axioms are

$$h \wedge f \to [r]f; \qquad h \wedge \neg f \to [r]\neg f; \qquad h \wedge d \to [r]d; \qquad h \wedge \neg d \to [r]\neg d.$$

Here, h is the *condition* under which the frame axioms are applicable.

In general, a positive frame axiom has the form

$$\mathit{FrCond}^+(\alpha, f) \wedge f \to [\alpha]f$$

and a negative frame axiom has the form

$$\mathit{FrCond}^-(\alpha, f) \wedge \neg f \to [\alpha]\neg f,$$

where $FrCond^+(\alpha, f)$ is a formula stating the conditions under which literal f remains positive and $FrCond^-(\alpha, f)$ is a formula stating the conditions under which literal $\neg f$ remains negative due to the execution of action α.

Instead of stating frame axioms directly, we shall use a slightly more concise expression by collecting all fluents invariant under the same conditions. We define the following abbreviation.

$$Inv(\alpha, \phi, \{f_1, \dots, f_m\}) \overset{def}{=} \text{ for } i \in \{1, \dots, m\},\ \phi \to ((f_i \to [\alpha]f_i) \wedge (\neg f_i \to [\alpha]\neg f_i)),$$

where $f_1, \dots, f_m \in \mathcal{F}$. $Inv(\alpha, \phi, \{f_1, \dots, f_m\})$ is called the *invariance predicate*. It is read 'When α is executed under condition ϕ, the truth values of fluents $f_1, \dots, f_m$ are invariant.

To relate frame axioms and invariance predicates, note that the following two statements hold ($\Rightarrow$ is read 'implies').

$$\mathcal{S}, w \models Inv(\alpha, FrCond^+(\alpha, f) \wedge f, F) \text{ s.t. } f \in F \Rightarrow \mathcal{S}, w \models FrCond^+(\alpha, f) \wedge f \to [\alpha]f.$$

$$\mathcal{S}, w \models Inv(\alpha, FrCond^-(\alpha, f) \wedge \neg f, F) \text{ s.t. } f \in F \Rightarrow \mathcal{S}, w \models FrCond^-(\alpha, f) \wedge \neg f \to [\alpha]\neg f.$$

Note the subtlety that the literal of the right polarity must be included in the condition of the invariance predicate.

We shall collect all invariance predicates in the set INV. Our approach assumes that for every/any α, for all $Inv(\alpha, \phi, F), Inv(\alpha, \phi', F') \in INV$, $\phi \wedge \phi' \equiv \perp$. Furthermore, for every effect axiom $\phi \to \Phi$ for α, for all $Inv(\alpha, \phi', F) \in INV$, either $\phi \wedge \phi' \equiv \perp$ or $\phi \models \phi'$. These assumptions keep things organized.

Now suppose we have the following invariance predicates (denoted INV_1).

$$\begin{array}{lll} Inv(g, f \wedge \neg h, \{d\}); & Inv(g, \neg f \wedge \neg h, \{f, d\}); & Inv(di, \neg f \wedge \neg d \wedge h, \{f, h\}); \\ Inv(r, h, \{f, d\}); & Inv(di, f \wedge \neg d \wedge h, \{h\}); & Inv(di, \neg f \wedge d \wedge h, \{f, d, h\}). \end{array}$$

INV_1 is a partial specification of action effects of the oil-drinking scenario. To further specify effects, one can supply the following effect axioms (denoted as PES_2).

$$\begin{array}{rcl} f \wedge \neg h & \to & [g]_{0.7}(f \wedge h) \wedge [g]_{0.3}\neg h; \\ \neg f \wedge \neg h & \to & [g]_{0.9}h \wedge [g]_{0.1}\neg h; \\ f \wedge \neg d \wedge h & \to & [di]_{0.85}(\neg f \wedge d) \wedge [di]_{0.15}(\neg f \wedge \neg d); \\ h & \to & [r]\neg h. \end{array}$$

Note that $\bigwedge_{\beta \in PES_1} \beta \equiv \bigwedge_{\delta \in INV_1 \cup PES_2} \delta$, but $INV_1 \cup PES_2$ is significantly smaller than PES_1.

Furthermore, we shall assume that effect and invariance specifications are complete, that is, that the knowledge engineer makes the completeness assumption about these specifications. Recall that $\langle\alpha\rangle\top$ abbreviates $\neg[\alpha]_0\top$ and note that $\mathcal{S}, w \models [\alpha]_q\varphi$ for $q > 0$ and $\varphi \not\equiv \bot$ if and only if $\mathcal{S}, w \models \neg[\alpha]_0\top$ (for all $\mathcal{S}$ and w). Therefore, if there is an effect axiom invariance predicate with condition ϕ, then one can assume the presence of an executability axiom $\phi \rightarrow \langle\alpha\rangle\top$. However, we must still specify that an action is inexecutable when none of the effect axiom conditions or invariance predicate conditions is met. Hence, the following *inexecutability* axiom is assumed present.[1]

$$\langle\alpha\rangle\top \rightarrow (\phi_1 \vee \cdots \vee \phi_j) \vee \bigvee_{\phi \in Cond^{Inv}(\alpha)} \phi,$$

where $\phi_1, \ldots, \phi_j$ are the conditions of the effect axioms for α and $Cond^{Inv}(\alpha)$ is the set of all the conditions mentioned in the invariance predicates for α. The inexecutability axioms for our example are

$$\langle g\rangle\top \rightarrow \neg h; \qquad \langle di\rangle\top \rightarrow h \wedge \neg(f \wedge d); \qquad \langle r\rangle\top \rightarrow h. \tag{1}$$

We shall collect all inexecutability axioms in the set INX. We shall refer to the set (1) in particular as INX_1.

5 From Underspecified to Completely Specified Transition Models

In this section, we define a *complete* transition model specification to capture the notion of a model which provides information about transitions between *every* pair of states. The two subsections provide approaches for completing partial specifications. A transition model specification is partial when some fluent in the specification is *underspecified*. The notion of an "underspecified" fluent is formalized by Definitions 5.2 and 5.3 below. We propose two alternative approaches for completing partial specifications in the subsections below. In the Section 6, it is proved that these approaches do in fact generate complete transition model specifications.

Definition 5.1 *A* complete transition model specification *(short: CTS) for action α is any set $B \subset \mathcal{L}$, such that, given C induced from $\mathcal{F}$, there exists a structure*

[1] Inexecutability axioms are also called *condition closure* axioms.

$\mathcal{S} = \langle C, R\rangle$ such that R is characterized by B. In other words, if there exists a structure $\mathcal{S} = \langle C, R\rangle$ where $(\alpha, R_\alpha) \in R$ and $\mathcal{S} \models \bigwedge_{\beta\in B} \beta$, and there is no $\mathcal{S}' = \langle C, R'\rangle$ such that $\mathcal{S}' \models \bigwedge_{\beta\in B} \beta$, where $(\alpha, R'_\alpha) \in R'$ and $R'_\alpha \neq R_\alpha$, then the set of sentences B is a complete transition model specification *for α.*

A PES is, in general, not a CTS: Let $\mathcal{F} = \{p_1\}$, $\alpha_1 \in \mathcal{A}$ and $B = \{p_1 \rightarrow [\alpha_1]p_1\}$. Then B is a PES for α_1. And let $w_1 \models p_1$ and $w_2 \models \neg p_1$. Assume $\mathcal{S}' \models p_1 \rightarrow [\alpha_1]p_1$, where $\mathcal{S}' = \langle C, R'\rangle$, $(w_2, w_2, 0.4) \in R'_{\alpha_1}$ and $(\alpha_1, R'_{\alpha_1}) \in R'$ and assume $\mathcal{S}'' \models p_1 \rightarrow [\alpha_1]p_1$, where $\mathcal{S}'' = \langle C, R''\rangle$, $(w_2, w_2, 0.5) \in R''_{\alpha_1}$ and $(\alpha_1, R''_{\alpha_1}) \in R''$. But the two structures $\mathcal{S}'$ and $\mathcal{S}''$ are different. Therefore, $p_1 \rightarrow [\alpha_1]p_1$ does not uniquely specify the accessibility relation for α_1. But the definition of a CTS says it must be unique.

Suppose a *completeness assumption* about effect axioms is as follows: The conditions of effect axioms for action α specifies all the conditions under which α has an effect, that is, under which α causes a fluent to change (see, e.g., Reiter Reiter [1991], §2.3). In deterministic systems, if one makes the completeness assumption about effect conditions, one can deduce frame axioms from the effect axioms Reiter [1991]. But effect axioms for non-deterministic systems are different, and frame axioms are not enough: Let $BK^{od} := INV_1 \cup PES_2 \cup INX_1$. Note that $BK^{od} \not\models f \wedge \neg h \rightarrow [g]_q(f \wedge d \wedge \neg h) \vee [g]_{q'}(\neg f \wedge d \wedge \neg h)$ for any q and q'. One could assume, due to lack of knowledge, that the truth value of f does not change, that is

$$BK^{od} \models f \wedge \neg h \rightarrow [g]_{0.3}(f \wedge d \wedge \neg h),$$

or one could assume a uniform distribution of probability over the possible values of f, that is

$$BK^{od} \models f \wedge \neg h \rightarrow [g]_{0.15}(f \wedge d \wedge \neg h) \wedge [g]_{0.15}(\neg f \wedge d \wedge \neg h).$$

There seems to be no clear way to decide between the two assumptions without knowledge of the domain; it depends on the domain of interest. In the following approaches dealing with this issue, we require two rather complicated definitions. Each definition is followed by an intuitive explanation.

Definition 5.2 *Let $\Upsilon^{\mathit{eff}}(\alpha, \phi) = \{\phi' \mid Inv(\alpha, \phi', F) \in INV, \phi \wedge \phi' \not\equiv \bot\}$. Given effect axiom $\phi \rightarrow [\alpha]_{q_1}\varphi_1 \wedge [\alpha]_{q_2}\varphi_2 \wedge \ldots \wedge [\alpha]_{q_n}\varphi_n$ for α of a PES, fluent f is* effectively underspecified *in effect $\varphi \in \{\varphi_1, \varphi_2, \cdots, \varphi_n\}$ under condition $\phi \wedge \phi'$ if $[\alpha]_q\varphi \not\models [\alpha]_q(\varphi \wedge f)$ and $[\alpha]_q\varphi \not\models [\alpha]_q(\varphi \wedge \neg f)$ and $f \notin F$, where $Inv(\alpha, \phi', F) \in INV$ for $\phi' \in \Upsilon^{\mathit{eff}}(\alpha, \phi)$.*

Intuitively, a fluent is underspecified in an effect φ (of the effect axiom for α with condition ϕ) under condition $\phi \wedge \phi'$, if its truth value is unknown after transition $[\alpha]_q\varphi$, and the invariance predicate for α with condition ϕ' does not reveal the truth value of the fluent.

Definition 5.3 *Let* $\Upsilon^{inv}(\alpha, \phi') = \{\phi \mid \phi \rightarrow \Phi$ *is an effect axiom for* $\alpha, \phi' \wedge \phi \not\equiv \bot\}$ *Given invariance predicate* $Inv(\alpha, \phi, F) \in INV$*, fluent* f *is* invariantly underspecified *under condition* $\phi' \wedge \neg \bigvee_{\phi \in Cond^{Inv}(\alpha)} \phi$ *if this condition is not a contradiction and if* $f \notin F$.

Intuitively, a fluent is underspecified under satisfiable condition $\phi' \wedge \neg \bigvee_{\phi \in Cond^{Inv}(\alpha)} \phi$, if the invariance predicate for α with condition ϕ' does not reveal the truth value of the fluent.

The definitions assume that all relevant information about effects of actions is contained in a clearly defined PES and set INV. If effect information were not easily located in this manner, it would be very difficult to 'complete' the specifications of effects of action as is done subsequently. In other words, our proposal for the management of probabilistic transition models includes the requirement that a PES and a set INV are clearly defined and accessible by the system or system-user.

When we say a fluent is underspecified, we mean it in the sense of one or more of Definitions 5.2 and 5.3. Next, we propose two alternative approaches for completing incompletely specified transition models. The approaches are: When a fluent is underspecified under a particular condition, (1) assume that it is invariant under that condition or (2) assume that it is uniformly distributed under that condition.

5.1 Always Assuming Invariance

For every $\alpha \in \mathcal{A}$ and $f \in \mathcal{F}$, for all conditions ϕ of effect axioms for α and conditions ϕ' of invariance predicates for α, if f is effectively underspecified in effect φ under condition $\phi \wedge \phi'$, assume the presence of *frame axioms*

$$\phi \wedge \phi' \wedge f \rightarrow [\alpha]_q(\varphi \wedge f) \text{ and } \phi \wedge \phi' \wedge \neg f \rightarrow [\alpha]_q(\varphi \wedge \neg f).$$

For every $\alpha \in \mathcal{A}$ and $f \in \mathcal{F}$, for all conditions ϕ' of invariance predicates for α, if f is invariantly underspecified under condition $\phi' \wedge \neg \bigvee_{\phi \in Cond^{Inv}(\alpha)} \phi$, assume the presence of frame axioms

$$\phi' \wedge \neg \bigvee_{\phi \in Cond^{Inv}(\alpha)} \phi \wedge f \rightarrow [\alpha]_q(\varphi \wedge f) \text{ and } \phi' \wedge \neg \bigvee_{\phi \in Cond^{Inv}(\alpha)} \phi \wedge \neg f \rightarrow [\alpha]_q(\varphi \wedge \neg f).$$

Given PES_1, INV is empty and there are thus no invariably underspecified fluents. However, f and d are effectively underspecified, leading to the assumption of these invariance formulae:

$$\begin{array}{rcl} f \wedge d \wedge \neg h & \rightarrow & [g]_{0.3}(d \wedge \neg h \wedge f); \\ f \wedge \neg d \wedge \neg h & \rightarrow & [g]_{0.3}(\neg d \wedge \neg h \wedge f); \\ \neg f \wedge \neg d \wedge h & \rightarrow & [di](\neg f \wedge \neg d \wedge h). \end{array} \quad (2)$$

Given PES_2 and INV_1: f is effectively underspecified in effect $\neg h$ under condition $(f \wedge \neg h) \wedge (f \wedge \neg h)$. The invariance formulae due to f are thus

$$(f \wedge \neg h) \wedge f \rightarrow [g]_{0.3}(\neg h \wedge f) \text{ and } (f \wedge \neg h) \wedge \neg f \rightarrow [g]_{0.3}(\neg h \wedge \neg f),$$

which simplifies to

$$f \wedge \neg h \rightarrow [g]_{0.3}(\neg h \wedge f). \quad (3)$$

d is invariantly underspecified under condition $(\neg f \wedge \neg d \wedge h) \wedge \neg(f \wedge \neg d \wedge h)$ (which is semantically equivalent to $\neg f \wedge \neg d \wedge h$). The invariance formulae due to d are thus

$$(\neg f \wedge \neg d \wedge h) \wedge d \rightarrow [di]d \text{ and } (\neg f \wedge \neg d \wedge h) \wedge \neg d \rightarrow [di]\neg d,$$

which simplifies to

$$\neg f \wedge \neg d \wedge h \rightarrow [di]\neg d. \quad (4)$$

Proposition 5.1 *$PES_1 \cup \{(2)\}$ is semantically equivalent to $PES_2 \cup INV_1 \cup \{(3), (4)\}$.*

Proof:
Observe that

$$\begin{array}{l} (f \wedge d \wedge \neg h \rightarrow [g]_{0.3}(d \wedge \neg h \wedge f)) \ \wedge \ (f \wedge \neg d \wedge \neg h \rightarrow [g]_{0.3}(\neg d \wedge \neg h \wedge f)) \\ \equiv Inv(g, f \wedge \neg h, \{d\}) \wedge (f \wedge \neg h \rightarrow [g]_{0.3}(\neg h \wedge f)), \end{array}$$

where the formulae on the LHS are in $\{(2)\}$ and on the RHS, $Inv(g, f \wedge \neg h, \{d\}) \in INV_1$ and $f \wedge \neg h \rightarrow [g]_{0.3}(\neg h \wedge f)$ is (3). And observe that

$$\begin{array}{l} \neg f \wedge \neg d \wedge h \rightarrow [di](\neg f \wedge \neg d \wedge h) \\ \quad \equiv (\neg f \wedge \neg d \wedge h \rightarrow [di](\neg f \wedge h)) \wedge (\neg f \wedge \neg d \wedge h \rightarrow [di]\neg d), \end{array}$$

where the formula on the LHS is in $\{(2)\}$ and on the RHS, $\neg f \wedge \neg d \wedge h \rightarrow [di](\neg f \wedge h) \in PES_1$ and $\neg f \wedge \neg d \wedge h \rightarrow [di]\neg d$ is (4). ■

Proposition 5.1 is supporting evidence that our strategies are correct.

5.2 Always Assuming Uniform Distribution

Let $U^{\mathit{eff}}(\alpha, \phi, \phi', \varphi) = \{f \in \mathcal{F} \mid f$ is effectively underspecified for α in effect φ under condition $\phi \wedge \phi'\}$, where ϕ is the condition of some effect axiom for α and ϕ' is the condition of some invariance predicate for α, and $U^{inv}(\alpha, \phi') = \{f \in \mathcal{F} \mid f$ is invariantly underspecified for α under condition $\phi' \wedge \neg \bigvee_{\phi \in Cond^{Inv}(\alpha)} \phi$ as described in Definition 5.3$\}$.

For every $\alpha \in \mathcal{A}$, for all conditions ϕ of effect axioms for α, for every transition $[\alpha]_q\varphi$ of the axiom, for all conditions ϕ' of invariance predicates for α, assume the presence of *equiprob formula*

$$\phi \rightarrow [\alpha]_{q_1}(\varphi \wedge \gamma_1) \wedge \cdots \wedge [\alpha]_{q_m}(\varphi \wedge \gamma_m),$$

where $\{\gamma_1, \ldots, \gamma_m\}$ are the $m = 2^{|U^{\mathit{eff}}(\alpha,\phi,\phi',\varphi)|}$ permutations of conjunctions of literals, given all the fluents in $U^{\mathit{eff}}(\alpha, \phi, \phi', \varphi)$ and $q_1 = \cdots = q_m = q/m$. For instance, if $U^{\mathit{eff}}(\alpha, \phi, \phi', \varphi) = \{f_2, f_4\}$ then the literal conjunction permutations are $\{f_2 \wedge f_4,\ f_2 \wedge \neg f_4,\ \neg f_2 \wedge f_4,\ \neg f_2 \wedge \neg f_4\}$.

For every action $\alpha \in \mathcal{A}$, for all conditions ϕ' of invariance predicate for α, assume the presence of *equiprob formula*

$$\phi' \wedge \neg \bigvee_{\phi \in Cond^{Inv}(\alpha)} \phi \rightarrow [\alpha]_{q_1}\gamma_1 \wedge \cdots \wedge [\alpha]_{q_n}\gamma_n,$$

where $\{\gamma_1, \ldots, \gamma_n\}$ are the $n = 2^{|U^{inv}(\alpha,\phi')|}$ permutations of conjunctions of literals, given all the fluents in $U^{inv}(\alpha, \phi')$ and $q_1 = \cdots = q_n = q/n$. Note that an equiprob formula need not be stated if its condition $\phi' \wedge \neg \bigvee_{\phi \in Cond^{Inv}(\alpha)} \phi$ is a contradiction. The same goes for cases when $U^{inv}(\alpha, \phi')$ is empty.

Given PES_1, INV is empty and there are thus no invariably underspecified fluents. However, f and d are effectively underspecified, leading to the assumption of these equiprob formulae:

$$\begin{array}{rcl} f \wedge d \wedge \neg h & \rightarrow & [g]_{0.15}(d \wedge \neg h \wedge f) \wedge [g]_{0.15}(d \wedge \neg h \wedge \neg f); \\ f \wedge \neg d \wedge \neg h & \rightarrow & [g]_{0.15}(\neg d \wedge \neg h \wedge f) \wedge [g]_{0.15}(\neg d \wedge \neg h \wedge \neg f); \\ \neg f \wedge \neg d \wedge h & \rightarrow & [d]_{0.5}(\neg f \wedge d \wedge h) \wedge [d]_{0.5}(\neg f \wedge \neg d \wedge h). \end{array} \tag{5}$$

Given INV_1 and PES_2: $U^{\mathit{eff}}(g, f \wedge \neg h, f \wedge \neg h, \neg h) = \{f\}$ (due to $[g]_{0.3}\neg h$ in effect axiom $f \wedge \neg h \rightarrow [g]_{0.7}(f \wedge h) \wedge [g]_{0.3}\neg h$). Hence, equiprob formula

$$f \wedge \neg h \rightarrow [g]_{0.15}(\neg h \wedge f) \wedge [g]_{0.15}(\neg h \wedge \neg f) \tag{6}$$

is assumed present. For all other conditions (combinations of ϕ and ϕ'), $U^{eff}(\alpha, \phi, \phi', \varphi) = \emptyset$. And due to invariance predicate $Inv(di, \neg f \wedge \neg d \wedge h, \{f, h\})$ in INV_1, $U^{inv}(di, \neg f \wedge \neg d \wedge h) = \{d\}$. Hence, equiprob formula

$$\neg f \wedge \neg d \wedge h \rightarrow [di]_{0.5} d \wedge [di]_{0.5} \neg d \tag{7}$$

is assumed present. For all other conditions ϕ', $U^{inv}(\alpha, \phi') = \emptyset$.

Proposition 5.2 $PES_1 \cup \{(5)\}$ *is semantically equivalent to* $PES_2 \cup INV_1 \cup \{(6), (7)\}$.

Proof:
Omitted. ■

6 The Two Approaches are Complete Specifications

This section presents a theorem (Thm. 6.1) which proves that there exists a process for transforming a PES, a set INV and the associated frame axioms or equiprob formulae into a complete transition model specification (CTS).

If there exists a world $w \in C$ such that $w \models \delta$, where δ is a propositional formula, and for all $w' \in C$, if $w' \neq w$ then $w' \not\models \delta$, we say that δ is *definitive* (then, δ defines a world; δ is a *complete propositional theory*). Let Def be the smallest set of all definitive formulae induced from $\mathcal{F}$.

Definition 6.1 *A* verbose effects specification *(VES) is a PES where all effect axiom conditions (the ϕ left of the $\rightarrow$) and effects (the φ right of the $\rightarrow$) are definitive formulae.*

Lemma 6.1 *Let INX^V be $\neg\langle\alpha\rangle\top \rightarrow \neg(\phi_1 \vee \cdots \vee \phi_j)$, where $\phi_1, \ldots, \phi_j$ are the j conditions of the j effect axioms in a VES V for α, and V is generated using the process described in the proof of Theorem 6.1. Then $INX^V \wedge \bigwedge_{\beta \in V} \beta$ is a CTS.*

Proof:
We must show that there exists a unique $R_\alpha : (C \times C) \mapsto \mathbb{Q}_{[0,1]}$ which is a total function from pairs of worlds into the rationals, and for every $w^- \in C$, either $\sum_{(w^-, w^+, pr) \in R_\alpha} pr = 1$ or $\sum_{(w^-, w^+, pr) \in R_\alpha} pr = 0$, such that $(\alpha, R_\alpha) \in R$ and $\langle C, R\rangle \models INX^V \wedge \bigwedge_{\beta \in V} \beta$.

For the sake of reference, let

$$\phi \to [\alpha]_{q_1}\varphi_1 \wedge [\alpha]_{q_2}\varphi_2 \wedge \ldots \wedge [\alpha]_{q_n}\varphi_n$$

be an arbitrary effect axiom of V. We may refer to the axiom as η. Construct R_α as follows: For all $w^-, w^+ \in C$: If $w^- \not\models (\phi_1 \vee \cdots \vee \phi_j)$, then $(w^-, w^+, 0) \in R_\alpha$. Else if $w^- \models \phi$: if $w^+ \models \varphi_k$ then $(w^-, w^+, q_k) \in R_\alpha$, else if $w^+ \not\models \varphi_1 \vee \cdots \vee \varphi_n$ then $(w^-, w^+, 0) \in R_\alpha$.

Now, the domain and co-domain of R_α are clearly adhered to. R_α is a *function* because of the constraint of a PES that for every i, for any pair of effects φ_{ik} and $\varphi_{ik'}$, $\varphi_{ik} \wedge \varphi_{ik'} \equiv \bot$, that is, never is more than one probability specified for reaching a world w^+ from some world w^-.

R_α is a *total* function because, given any pair $(w^-, w^+) \in (C \times C)$, if $w^- \models \psi_i$ where ϕ_i is the condition of the i-th effect axiom, then either (i) $w^+ \models \varphi_{ik}$ for some transition $[\alpha]_q\varphi_{ik}$ in the axiom, in which case $(w^-, w^+, q) \in R_\alpha$ or (ii) $w^+ \not\models \varphi_{ik}$ for all transitions in the axiom, in which case $(w^-, w^+, 0) \in R_\alpha$, due to the PES constraint that the transition probabilities $q_{i1}, \ldots, q_{in}$ of any axiom i must sum to 1. Else, for all $w^- \in C$ such that $w^- \models \neg(\phi_1 \vee \cdots \vee \phi_j)$, $(w^-, w^+, 0) \in R_\alpha$ for all $w^+ \in C$, due to the PES constraint that for any pair of conditions ϕ_i and $\phi_{i'}$, $\phi_i \wedge \phi_{i'} \equiv \bot$. It follows implicitly that for every $w^- \in C$, either $\sum_{(w^-, w^+, pr) \in R_\alpha} pr = 1$ or $\sum_{(w^-, w^+, pr) \in R_\alpha} pr = 0$.

Simply, by construction of R_α, it follows that $\langle C, R \rangle \models INX^V$. And as a direct consequence of the construction of R_α, it follows that $\langle C, R \rangle \models \bigwedge_{\beta \in V} \beta$.

We shall now show that no other R'_α ($\neq R_\alpha$) can be constructed such that $(\alpha, R'_\alpha) \in R'$ and $\langle C, R' \rangle \models INX^V \wedge \bigwedge_{\beta \in V} \beta$. Let (w^-, w^+, q_k) be some element of R_α as constructed. Let $q' \in \mathbb{Q}_{[0,1]}$ such that $|q' - q_k| > 0$. If $w^- \not\models (\phi_1 \vee \cdots \vee \phi_j)$, then $(w^-, w^+, q') \in R'_\alpha$, where $q' > 0$. But then $\langle C, R' \rangle \not\models INX^V$. And if $w^- \models \phi$ and $w^+ \models \varphi_k$, then $q_k \neq q'$ and $\langle C, R' \rangle \not\models \phi \to [\alpha]_{q_k}\varphi_k$, which implies that $\langle C, R' \rangle \not\models \eta$, which implies that $\langle C, R' \rangle \not\models \bigwedge_{\beta \in V} \beta$. Else, if $w^- \models \phi$ and $w^+ \not\models \varphi_1 \vee \cdots \vee \varphi_n$ then $(w^-, w^+, q') \in R_\alpha$ where $q' \neq 0$. But this is a contradiction, because it is required that $\sum_{(w^-, w^+, pr) \in R_\alpha} pr = 1$, but due to the PES constraint that the transition probabilities $q_{i1}, \ldots, q_{in}$ of any axiom i must sum to 1, $\sum_{(w^-, w^+, pr) \in R_\alpha} pr > 1$. ∎

Definition 6.2 *Note that any formula $\phi \to \Phi$ (with condition ϕ) such that $\phi \not\models f$ and $\phi \not\models \neg f$ for some $f \in \mathcal{F}$, is semantically equivalent to $(\phi \wedge f \to \Phi) \wedge (\phi \wedge \neg f \to \Phi)$. Given this observation, one can expand any formula of the form $\phi \to \Phi$ into a set of semantically equivalent formulae, each with a definitive condition.*

We shall refer to this process applied to a set X of formulae as expending X by conditions.

Proposition 6.1 $(\phi \rightarrow [\alpha]\varphi) \wedge (\phi' \rightarrow [\alpha]_q\varphi') \wedge (\phi' \rightarrow \phi) \models \phi' \rightarrow [\alpha]_q(\varphi \wedge \varphi')$ *for all* $q \in \mathbb{Q}_{[0,1]}$.

Proof:
Let $\mathcal{S}$ be an arbitrary SLAP structure and w a world in it. Suppose $\mathcal{S}, w \models (\phi \rightarrow [\alpha]\varphi) \wedge (\phi' \rightarrow [\alpha]_q\varphi') \wedge (\phi' \rightarrow \phi)$. Assume $\mathcal{S}, w \models \phi'$. Then $\mathcal{S}, w \models \phi$ (and thus $\mathcal{S}, w \models [\alpha]\varphi$) and $\mathcal{S}, w \models [\alpha]_q\varphi'$ (i.e., $\mathcal{S}, w \models \phi \wedge [\alpha]\varphi \wedge [\alpha]_q\varphi'$.

Now, $\mathcal{S}, w \models [\alpha]\varphi \wedge [\alpha]_q\varphi'$ iff

$$\sum_{(w,w',pr) \in R_\alpha, \mathcal{S}, w' \models \varphi} pr = 1 \text{ and } \sum_{(w,w',pr) \in R_\alpha, \mathcal{S}, w' \models \varphi'} pr = q.$$

Therefore, $\{(w, w', pr) \in R_\alpha \mid w' \models \varphi'\} \subseteq \{(w, w', pr) \in R_\alpha \mid w' \models \varphi, pr > 0\}$. Hence, $\mathcal{S}, w \models$ iff $\sum_{(w,w',pr) \in R_\alpha, \mathcal{S}, w' \models \varphi \wedge \varphi'} pr = q$, and by definition, $\mathcal{S}, w \models [\alpha]_q(\varphi \wedge \varphi')$. ∎

Proposition 6.2 *A fluent cannot be effectively and invariantly underspecified under the same condition.*

Proof:
By definition of $\Upsilon^{inv}(\alpha, \phi')$, whenever $\phi \wedge \phi' \not\equiv \bot$, then $\phi \in \Upsilon^{inv}(\alpha, \phi')$. Note that $\neg \bigvee_{\phi \in Cond^{Inv}(\alpha)} \phi \equiv \neg\phi_1 \wedge \neg\phi_2 \wedge \cdots \wedge \neg\phi_n$, where $\Upsilon^{inv}(\alpha, \phi') = \{\phi_1, \phi_2, \ldots, \phi_n\}$. But $\phi \in \{\phi_1, \phi_2, \ldots, \phi_n\}$. Therefore, $\phi \wedge \neg \bigvee_{\phi \in Cond^{Inv}(\alpha)} \phi \equiv \bot$ and the proposition follows. ∎

Theorem 6.1 *For both approaches, given a PES* Π *for* α*, a set of invariance predicates* INV *for* α*, an inexecutability axioms* INX *for* α *derived from* Π *and* INV*, a set of frame axioms* FA *for* α *and a set of equiprob formulae* EF *for* α*, their union is a CTS for* α.

Proof:
Suppose V is a VES for α and INX^V is an inexecutability axiom derived from V as in Lemma 6.1. If we can show that V and INX^V exist such that $INX \wedge \bigwedge_{\beta \in \Pi \cup INV \cup FA \cup EF} \equiv INX^V \wedge \bigwedge_{\delta \in V} \delta$, then by Lemma 6.1, we have proved the theorem. Hence, we show how to convert $\Pi \cup INV \cup FA \cup EF$ into a semantically equivalent VES V and we prove that $INX \equiv INX^V$.

We show how to 'enlarge' Π into a VES using four rewrite rules involving INV, IF and EF. The rewrite rules are:

1. For every $Inv(\alpha, \phi, \{f_1, \ldots, f_m\}) \in INV$, for $i \in \{1, \ldots, m\}$, add $\phi \wedge f_i \rightarrow [\alpha] f_i$ and $\phi \wedge \neg f_i \rightarrow [\alpha] \neg f_i$ to INV if $\phi \wedge f_i$, respectively, $\phi \wedge \neg f_i$ is satisfiable. Remove all invariance predicates from INV.

2. Expand Π, INV, FF and EF by conditions.

3. For every $\phi' \rightarrow [\alpha]\ell \in INV$, if there is no $\phi \rightarrow \Phi \in \Pi$ such that $\phi \equiv \phi'$, then add $\phi' \rightarrow [\alpha]\top$ to Π.

4. When assuming invariance: Note that EF is empty. For every $\phi \rightarrow \Phi \in \Pi$, for every $\phi' \rightarrow [\alpha]_q \varphi' \in FA$ such that $\phi \equiv \phi'$, for every $[\alpha]_q \varphi$ in Φ, if $\varphi' \models \varphi$, then replace $[\alpha]_q \varphi$ by $[\alpha]_q \varphi'$.

   When assuming uniform distribution: Note that FA is empty. For every $\phi \rightarrow \Phi \in \Pi$, for every $\phi \rightarrow [\alpha]_{q_1}(\varphi \wedge \gamma_1) \wedge \cdots \wedge [\alpha]_{q_m}(\varphi \wedge \gamma_m), \in EF$ such that $\phi \equiv \phi'$, for every $[\alpha]_q \varphi$ in Φ, if $\varphi_i \models \varphi$ for all $i \in \{1, \ldots, m\}$, then replace $[\alpha]_q \varphi$ by $[\alpha]_{q_1}(\varphi \wedge \gamma_1) \wedge \cdots \wedge [\alpha]_{q_m}(\varphi \wedge \gamma_m)$.

5. For every $\phi \rightarrow \Phi \in \Pi$, if $\phi' \rightarrow [\alpha]\ell \in INV$ such that $\phi \equiv \phi'$, then replace every $[\alpha]\varphi$ in Φ by $[\alpha](\varphi \wedge \ell)$.

Recalling that $Inv(\alpha, \phi, \{f_1, \ldots, f_m\})$ abbreviates: for $i \in \{1, \ldots, m\}$, $\phi \rightarrow ((f_i \rightarrow [\alpha] f_i) \wedge (\neg f_i \rightarrow [\alpha] \neg f_i))$, step 1 is sound. By Definition 6.2, step 2 is sound. Due to $\phi' \rightarrow [\alpha]\ell \models \phi' \rightarrow [\alpha]\top$, step 3 is sound. Note that, by Proposition 6.2, no formula in Π is modified/rewritten twice in step 4. Hence, the process in step 4 does not cause side-effects. And together with the two assumptions and the respective explanations of the approaches to complete the specifications, step 4 is sound. By Proposition 6.1, step 5 is sound.

Note that the two approaches of Sections 5.1 and 5.2 are designed exactly to deal with fluents not dealt with before, that is, to appropriately adds literals corresponding to every fluent not mentioned in an effect of some effect axiom, for every effect in every axiom. Thus, by the nature of INV, FA and EF and the rewrite rules of steps 4 and 5, every effect of every effect axiom is now a definitive formula.

A VES is a PES; the conditions of all its effect axioms must thus be disjoint. It is assumed that Π is initially a PES, hence, with disjoint conditions. None of the rewrite rules causes some pair of conditions in Π to be joint: In particular, step 2: Expansion by conditions cannot cause joint conditions.

Observe that INX depends only on the axiom conditions of the original Π, which has essentially the same axiom conditions as those of V (given that condition expansion in step 2 does not add conditions), and INX^V depends only on the axiom conditions of V. Hence $INX \equiv INX^V$. ∎

6.1 Example

PES_2 and INV_1 presented in Section 4 are repeated here, for convenience:

INV_1:

$$\begin{aligned}
&Inv(g, f \wedge \neg h, \{d\});\\
&Inv(g, \neg f \wedge \neg h, \{f, d\});\\
&Inv(di, \neg f \wedge \neg d \wedge h, \{f, h\});\\
&Inv(di, \neg f \wedge d \wedge h, \{f, d, h\});\\
&Inv(di, f \wedge \neg d \wedge h, \{h\});\\
&Inv(r, h, \{f, d\}).
\end{aligned}$$

PES_2:

$$\begin{aligned}
f \wedge \neg h &\rightarrow [g]_{0.7}(f \wedge h) \wedge [g]_{0.3}\neg h;\\
\neg f \wedge \neg h &\rightarrow [g]_{0.9}h \wedge [g]_{0.1}\neg h;\\
f \wedge \neg d \wedge h &\rightarrow [di]_{0.85}(\neg f \wedge d) \wedge [di]_{0.15}(\neg f \wedge \neg d);\\
h &\rightarrow [r]\neg h.
\end{aligned}$$

When assuming invariance (AI) by default, FA:

$$\begin{aligned}
f \wedge \neg h &\rightarrow [g]_{0.3}(\neg h \wedge f);\\
\neg f \wedge \neg d \wedge h &\rightarrow [di]\neg d.
\end{aligned}$$

When assuming uniform distribution (AUD) by default, EF:

$$\begin{aligned}
f \wedge \neg h &\rightarrow [g]_{0.15}(\neg h \wedge f) \wedge [g]_{0.15}(\neg h \wedge \neg f);\\
\neg f \wedge \neg d \wedge h &\rightarrow [di]_{0.5}d \wedge [di]_{0.5}\neg d.
\end{aligned}$$

The transformation process is now applied to these sentences, using the five rewrite rules presented in the proof of Theorem 6.1.

Rule 1. INV_1:

$f \wedge \neg h \wedge d \rightarrow [g]d$; $\quad$ $f \wedge \neg h \wedge \neg d \rightarrow [g]\neg d$;
$\neg f \wedge \neg h \rightarrow [g]\neg f$; $\quad$ $\neg f \wedge \neg h \wedge d \rightarrow [g]d$; $\quad$ $\neg f \wedge \neg h \wedge \neg d \rightarrow [g]\neg d$;
$\neg f \wedge \neg d \wedge h \rightarrow [di]\neg f$; $\quad$ $\neg f \wedge \neg d \wedge h \rightarrow [di]h$;
$\neg f \wedge d \wedge h \rightarrow [di]\neg f$; $\quad$ $\neg f \wedge d \wedge h \rightarrow [di]d$; $\quad$ $\neg f \wedge d \wedge h \rightarrow [di]h$;
$f \wedge \neg d \wedge h \rightarrow [di]h$;
$h \wedge f \rightarrow [r]f$; $\quad$ $h \wedge \neg f \rightarrow [r]\neg f$; $\quad$ $h \wedge d \rightarrow [r]d$; $\quad$ $h \wedge \neg d \rightarrow [r]\neg d$.

Rule 2. PES_2:

$$\begin{array}{rcl}
f \wedge \neg h \wedge d & \rightarrow & [g]_{0.7}(f \wedge h) \wedge [g]_{0.3}\neg h; \\
f \wedge \neg h \wedge \neg d & \rightarrow & [g]_{0.7}(f \wedge h) \wedge [g]_{0.3}\neg h; \\
\neg f \wedge \neg h \wedge d & \rightarrow & [g]_{0.9}h \wedge [g]_{0.1}\neg h; \\
\neg f \wedge \neg h \wedge \neg d & \rightarrow & [g]_{0.9}h \wedge [g]_{0.1}\neg h; \\
f \wedge \neg d \wedge h & \rightarrow & [di]_{0.85}(\neg f \wedge d) \wedge [di]_{0.15}(\neg f \wedge \neg d); \\
h \wedge f \wedge d & \rightarrow & [r]\neg h; \\
h \wedge f \wedge \neg d & \rightarrow & [r]\neg h; \\
h \wedge \neg f \wedge d & \rightarrow & [r]\neg h; \\
h \wedge \neg f \wedge \neg d & \rightarrow & [r]\neg h.
\end{array}$$

When AI, FA:

$$\begin{array}{rcl}
f \wedge \neg h \wedge d & \rightarrow & [g]_{0.3}(\neg h \wedge f); \\
f \wedge \neg h \wedge \neg d & \rightarrow & [g]_{0.3}(\neg h \wedge f); \\
\neg f \wedge \neg d \wedge h & \rightarrow & [di]\neg d.
\end{array}$$

When AUD, EF:

$$\begin{array}{rcl}
f \wedge \neg h \wedge d & \rightarrow & [g]_{0.15}(\neg h \wedge f) \wedge [g]_{0.15}(\neg h \wedge \neg f); \\
f \wedge \neg h \wedge \neg d & \rightarrow & [g]_{0.15}(\neg h \wedge f) \wedge [g]_{0.15}(\neg h \wedge \neg f); \\
\neg f \wedge \neg d \wedge h & \rightarrow & [di]_{0.5}d \wedge [di]_{0.5}\neg d.
\end{array}$$

INV_1:

$f \wedge \neg h \wedge d \rightarrow [g]d; \qquad f \wedge \neg h \wedge \neg d \rightarrow [g]\neg d;$
$\neg f \wedge \neg h \wedge d \rightarrow [g]\neg f; \qquad \neg f \wedge \neg h \wedge \neg d \rightarrow [g]\neg f;$
$\neg f \wedge \neg h \wedge d \rightarrow [g]d; \qquad \neg f \wedge \neg h \wedge \neg d \rightarrow [g]\neg d;$
$\neg f \wedge \neg d \wedge h \rightarrow [di]\neg f; \qquad \neg f \wedge \neg d \wedge h \rightarrow [di]h;$
$\neg f \wedge d \wedge h \rightarrow [di]\neg f; \qquad \neg f \wedge d \wedge h \rightarrow [di]d; \qquad \neg f \wedge d \wedge h \rightarrow [di]h;$
$f \wedge \neg d \wedge h \rightarrow [di]h;$
$h \wedge f \wedge d \rightarrow [r]f; \quad h \wedge \neg f \wedge d \rightarrow [r]\neg f; \quad h \wedge d \wedge f \rightarrow [r]d; \quad h \wedge \neg d \wedge f \rightarrow [r]\neg d;$
$h \wedge f \wedge \neg d \rightarrow [r]f; \quad h \wedge \neg f \wedge \neg d \rightarrow [r]\neg f; \quad h \wedge d \wedge \neg f \rightarrow [r]d; \quad h \wedge \neg d \wedge \neg f \rightarrow [r]\neg d.$

Rule 3. The following are added to PES_2.

$$\begin{aligned} \neg f \wedge \neg d \wedge h &\rightarrow [di]\top; \\ \neg f \wedge d \wedge h &\rightarrow [di]\top. \end{aligned}$$

Rule 4. When AI, the changes in PES_2 are:

$$\begin{aligned} f \wedge \neg h \wedge d &\rightarrow [g]_{0.7}(f \wedge h) \wedge [g]_{0.3}(\neg h \wedge f); \\ f \wedge \neg h \wedge \neg d &\rightarrow [g]_{0.7}(f \wedge h) \wedge [g]_{0.3}(\neg h \wedge f); \\ \neg f \wedge \neg d \wedge h &\rightarrow [di]\neg d. \end{aligned}$$

When AUD, the changes in PES_2 are:

$$\begin{aligned} f \wedge \neg h \wedge d &\rightarrow [g]_{0.7}(f \wedge h) \wedge [g]_{0.15}(\neg h \wedge f) \wedge [g]_{0.15}(\neg h \wedge \neg f); \\ f \wedge \neg h \wedge \neg d &\rightarrow [g]_{0.7}(f \wedge h) \wedge [g]_{0.15}(\neg h \wedge f) \wedge [g]_{0.15}(\neg h \wedge \neg f); \\ \neg f \wedge \neg d \wedge h &\rightarrow [di]_{0.5}d \wedge [di]_{0.5}\neg d. \end{aligned}$$

Rule 5. When AI, PES_2:

$$\begin{array}{rcl} f \wedge \neg h \wedge d & \to & [g]_{0.7}(f \wedge h) \wedge [g]_{0.3}(\neg h \wedge f); \\ f \wedge \neg h \wedge \neg d & \to & [g]_{0.7}(f \wedge h) \wedge [g]_{0.3}(\neg h \wedge f); \\ \neg f \wedge \neg h \wedge d & \to & [g]_{0.9}(h \wedge \neg f \wedge d) \wedge [g]_{0.1}(\neg h \wedge \neg f \wedge d); \\ \neg f \wedge \neg h \wedge \neg d & \to & [g]_{0.9}(h \wedge \neg f \wedge \neg d) \wedge [g]_{0.1}(\neg h \wedge \neg f \wedge \neg d); \\ f \wedge \neg d \wedge h & \to & [di]_{0.85}(\neg f \wedge d \wedge h) \wedge [di]_{0.15}(\neg f \wedge \neg d \wedge h); \\ \neg f \wedge \neg d \wedge h & \to & [di](\neg d \wedge \neg f \wedge h); \\ \neg f \wedge d \wedge h & \to & [di](\top \wedge \neg f \wedge d \wedge h); \\ h \wedge f \wedge d & \to & [r](\neg h \wedge f \wedge d); \\ h \wedge f \wedge \neg d & \to & [r](\neg h \wedge \neg d \wedge f); \\ h \wedge \neg f \wedge d & \to & [r](\neg h \wedge \neg f \wedge d); \\ h \wedge \neg f \wedge \neg d & \to & [r](\neg h \wedge \neg f \wedge \neg d). \end{array}$$

When AUD, PES_2:

$$\begin{array}{rcl} f \wedge \neg h \wedge d & \to & [g]_{0.7}(f \wedge h) \wedge [g]_{0.15}(\neg h \wedge f) \wedge [g]_{0.15}(\neg h \wedge \neg f); \\ f \wedge \neg h \wedge \neg d & \to & [g]_{0.7}(f \wedge h) \wedge [g]_{0.15}(\neg h \wedge f) \wedge [g]_{0.15}(\neg h \wedge \neg f); \\ \neg f \wedge \neg h \wedge d & \to & [g]_{0.9}(h \wedge \neg f \wedge d) \wedge [g]_{0.1}(\neg h \wedge \neg f \wedge d); \\ \neg f \wedge \neg h \wedge \neg d & \to & [g]_{0.9}(h \wedge \neg f \wedge \neg d) \wedge [g]_{0.1}(\neg h \wedge \neg f \wedge \neg d); \\ f \wedge \neg d \wedge h & \to & [di]_{0.85}(\neg f \wedge d \wedge h) \wedge [di]_{0.15}(\neg f \wedge \neg d \wedge h); \\ \neg f \wedge \neg d \wedge h & \to & [di]_{0.5}(d \wedge \neg f \wedge h) \wedge [di]_{0.5}(\neg d \wedge \neg f \wedge h); \\ \neg f \wedge d \wedge h & \to & [di](\top \wedge \neg f \wedge d \wedge h); \\ h \wedge f \wedge d & \to & [r](\neg h \wedge f \wedge d); \\ h \wedge f \wedge \neg d & \to & [r](\neg h \wedge \neg d \wedge f); \\ h \wedge \neg f \wedge d & \to & [r](\neg h \wedge \neg f \wedge d); \\ h \wedge \neg f \wedge \neg d & \to & [r](\neg h \wedge \neg f \wedge \neg d). \end{array}$$

The following is the set of inexecutability axioms INX_2 derived from the new PES_2:

$$\begin{array}{rcl} \langle g \rangle & \to & (f \wedge \neg h \wedge d) \vee (f \wedge \neg h \wedge \neg d) \vee (\neg f \wedge \neg h \wedge d) \vee (\neg f \wedge \neg h \wedge \neg d); \\ \langle di \rangle & \to & (f \wedge \neg d \wedge h) \vee (\neg f \wedge \neg d \wedge h) \vee (\neg f \wedge d \wedge h); \\ \langle r \rangle & \to & (h \wedge f \wedge d) \vee (h \wedge f \wedge \neg d) \vee (h \wedge \neg f \wedge d) \vee (h \wedge \neg f \wedge \neg d) \end{array}$$

which is semantically equivalent to

$$\begin{array}{rcl} \langle g \rangle \top & \rightarrow & \neg h; \\ \langle di \rangle \top & \rightarrow & h \wedge \neg(f \wedge d); \\ \langle r \rangle \top & \rightarrow & h. \end{array}$$

which is identical to INX_1 (Formulae 1, page 10). Furthermore, PES_2 is a VES, and it is thus easy to see that $PES_2 \cup INX_2$ is a set of three complete transition model specifications for the three actions.

7 Concluding Remarks

The work appearing in this report has been presented at a workshop Rens et al. [2013a]. However, Definitions 5.2 and 5.3 above, replace three definitions in that paper. With the new definitions, the two approaches corresponding to the two assumptions (Secs. 5.1 and 5.2) are slightly simpler. Furthermore, these two approaches for generating full specifications, presented in this report, do not constrain the knowledge engineer as much: In the approach presented in the workshop paper, the following constraints were in place. For every effect axiom $\phi \rightarrow \Phi$ for α, for all $Inv(\alpha, \phi', F) \in INV$, either $\phi \wedge \phi' \equiv \bot$ or $\phi \models \phi'$. These constraints are no longer necessary with the new approach.

We proved that our two approaches to complete underspecified transition models results in complete transition model specifications.

There seems to be two issues with underspecified models. One is knowing what information is missing. The other is deciding what information to add and how to add it correctly and completely. We have presented a systematic approach to generate complete specifications of probabilistic transition models (CTSs) with a probabilistic modal logic. For these specifications to be more compact than they would be if transition probabilities were simply written down, it is expected that a user/knowledge engineer will capture (with sentences of a logical language) some transition information from the domain of interest, and then for missing information, express the desired transition behavior of the model of the domain, and finally, for information still not provided by the user, s/he must take a stance as to what the default transition behavior should be: invariance of the truth values of fluents not mentioned in the effect axioms, or uniform distribution of transition probabilities. In real world situations, a combination of assumptions may be more effective. For instance, in a very dynamic environment, the default should perhaps be ‘variance’. That is, when information is not given about how the truth value of a fluent should change when some action is executed, it could be assumed that the fluent’s value will *always* change. Nevertheless, assuming (necessary) (in)variance is an assumption of certainty; these are ‘minimum entropy’/certain information assumptions and could be studied under the topic of traditional non-monotonic reasoning Brewka [2012].

The ‘uniform distribution’ assumption on the other hand is a kind of ‘maximum entropy’ approach. Wang and Schmolze Wang and Schmolze [2005] have a very similar approach to ours to achieve compact representations in POMDP planning. Some researchers (see, e.g., Grove et al. [1994] and the work of Kern-Isberner Kern-Isberner [2001] and colleagues) have proposed the assignment of a unique probability distribution over a vocabulary such that information theoretic entropy is maximized while the available probabilistic information is conserved. This *principle of maximum entropy* Jaynes [1978] seems to be a reasonable approach, but it may also be reasonable to assume a particular *a priori* probability distribution for a given domain when no other information is forthcoming. Although “default reasoning about probabilities” Jaeger [1994] is usually applied to what is believed in the *current* situation, the idea is easily applied to what will be believed in the *next* situation, that is, to transition models.

Another approach to more compact specifications is via notions of conditional independence of Belief Networks. See, for example, Fierens *et al.* Fierens et al. [2005] for a starting point in the area of combining belief nets with logic. We have not looked at the relationship between the notion of invariance and conditional independence in a probabilistic setting.

Bacchus, Halpern and Levesque Bacchus et al. [1999] give an account of specifying stochastic actions in the situation calculus while retaining Reiter’s solution to the frame problem Reiter [1991] via successor-state axioms (SSAs). In particular, §3 of their paper shows how to deal with a nondetrministic action by ‘decomposing’ it into a set of deterministic actions, each leading to one of the effects of the nondetrministic action. In SLAP, stochastic (nondeterministic) actions are specified ‘directly’, actions are not decomposed. The ‘direct approach’ corresponds more closely to POMDP models than the ‘decomposition approach’, and thus aligns better with logics with explicit POMDP semantics. We could thus not rely on Reiter’s solution. A deeper study is needed to compare the pros and cons of using decomposition and SSAs, on the one hand, and using our direct approach without SSAs, on the other hand.

References

F. Bacchus, J. Halpern, and H. Levesque. Reasoning about noisy sensors and effectors in the situation calculus. *Artificial Intelligence*, 111(1–2):171–208, 1999.

G. Brewka. *Nonmonotonic Reasoning: Logical Foundations of Commonsense (Cambridge Tracts in Theorl. Comp. Sci.)*. Cambridge University Press, reissue edition, 2012.

D. Fierens, H. Blockeel, M. Bruynooghe, and J. Ramon. Logical bayesian networks and their relation to other probabilistic logical models. In S. Kramer and B. Pfahringer, editors, *ILP 2005*, volume 3625 of *LNAI*, pages 121–135, Heidelberg/Berlin, 2005. Springer Verlag.

A. Grove, J. Halpern, and D. Koller. Random worlds and maximum entropy. *Journal of Artificial Intelligence Research (JAIR)*, 2:33–88, 1994.

M. Jaeger. A logic for default reasoning about probabilities. In *Proc. of 10th Intl. Conf. on Uncertainty in Artif. Intell.*, UAI'94, pages 352–359, San Francisco, CA, USA, 1994. Morgan Kaufmann Publishers Inc. ISBN 1-55860-332-8. URL `http://dl.acm.org/citation.cfm?id=2074394.2074439`.

E. Jaynes. Where do we stand on maximum entropy? In *The Maximum Entropy Formalism*, pages 15–118. MIT Press, 1978.

G. Kern-Isberner. Conditionals in nonmonotonic reasoning and belief revision. In *Lecture Notes in Computer Science (LNCS)*, number 2087. Springer Verlag, Berlin/Heidelberg, 2001.

J. McCarthy and P. Hayes. Some philosophical problems from the standpoint of artificial intelligence. *Machine Intelligence*, 4:463–502, 1969.

R. Reiter. The frame problem in the situation calculus: A simple solution (sometimes) and a completeness result for goal regression. In V. Lifschitz, editor, *Artificial intelligence and mathematical theory of computation: papers in honor of John McCarthy*, pages 359–380. Academic Press Professional, Inc., San Diego, CA, USA, 1991.

G. Rens, T. Meyer, and G. Lakemeyer. On the logical specification of probabilistic transition models. In *Proc. of 11th Intl. Symposium on Logical Formalizations of Commonsense Reasoning (COMMONSENSE 2013)*, May 2013a.

G. Rens, T. Meyer, and G. Lakemeyer. SLAP: Specification logic of actions with probability. *Journal of Applied Logic*, 2013b. ISSN 1570-8683. url = http://dx.doi.org/10.1016/j.jal.2013.09.001.

C. Wang and J. Schmolze. Planning with POMDPs using a compact, logic-based representation. In *Proc. of 17th IEEE Intl. Conf. on Tools with Artif. Intell. (ICTAI'05)*, pages 523–530, Los Alamitos, CA, USA, 2005. IEEE Computer Society.