

On the Logical Specification of Probabilistic Transition Models

Gavin Rens¹ and Thomas Meyer¹ and Gerhard Lakemeyer²

¹Centre for Artificial Intelligence Research
CSIR Meraka and Uni. of KwaZulu-Natal, South Africa

²RWTH Aachen University, Germany

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 probabilistic transitions due to the execution of actions of some agent. We propose a definition for ‘proper’ and ‘full’ 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.

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 paper, we investigate strategies for smaller ‘full’ specifications which follow intuitive lines of reasoning, 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 knows that she or he needs not (re)state the information.

The first author’s present research is in the development of a modal logic with a partially observable Markov decision process (POMDP) semantics. A sublogic is the Specification Logic of Actions with Probability (SLAP). Our investigation into the issues in specifying domains with SLAP lead to the work presented here. For us, it thus makes sense to tackle these issues using SLAP. However, the research discussed in this paper, may well be applied to other logics with probabilistic transition semantics.

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 (one cannot model the (epistemic) effects of observations with SLAP). The robot has the set of (intended) actions $\mathcal{A} = \{\text{grab}, \text{drink}, \text{weigh}, \text{replace}\}$ with expected meanings. The robot experiences its environment through three Boolean features: $\mathcal{P} = \{\text{full}, \text{drank}, \text{holding}\}$ 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 $(\neg\text{full} \wedge \neg\text{holding}) \rightarrow ([\text{grab}]_{0.9}(\neg\text{full} \wedge \text{holding}) \wedge [\text{grab}]_{0.1}(\neg\text{full} \wedge \neg\text{holding}))$ follow from BK ?

Next, we define SLAP. Then we investigate the issues in specifying probabilistic transition models and we propose some methods for their specification. Finally, our proposed methods are proved correct in a sense to be defined later.

Specification Logic of Actions with Probability

To illustrate our ideas, we present a modal logic for specifying agents’ stochastic action models. First we present the syntax, then the semantics of SLAP.

Syntax

The vocabulary of our language contains three sorts:

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

From now on, we denote $\mathbb{Q} \cap [0, 1]$ 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 1 Let $\alpha \in \mathcal{A}$, $q \in \mathbb{Q}_{[0,1]}$ and $p, p_1, \dots, p_m \in \mathcal{P}$. The language of SLAP, denoted \mathcal{L}_{SLAP} , is the least set of Ψ defined by the grammar:

$$\begin{aligned} \varphi &::= p \mid \top \mid \neg\varphi \mid \varphi \wedge \varphi. \\ \Phi &::= \varphi \mid \neg\Phi \mid \Phi \wedge \Phi \mid [\alpha]_q\varphi. \\ \Psi &::= \Box\Phi \mid Inv(\alpha, \varphi, \{p_1, \dots, p_m\}). \end{aligned}$$

Note that formulae with nested modal operators of the form $\Box\Box\Phi$, $\Box\Box\Box\Phi$, etc. or of the form $[\alpha]_q[\alpha]_q\varphi$, $[\alpha]_q[\alpha]_q[\alpha]_q\varphi$, etc. are not in \mathcal{L}_{SLAP} . ‘Single-step’ formulae are sufficient to specify action transitions. Formulae of the form $\Box[\alpha]_q\varphi$ are allowed. As usual, we treat \perp, \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.

$[\alpha]_q\varphi$ 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 ’. One reads $\Box\Phi$ as ‘ Φ holds in every possible world’. We require the \Box operator to mark certain sentences (axioms) which model the domain of interest as holding in *all* possible worlds. $Inv(\alpha, \psi, \{p_1, \dots, p_m\})$ is called the *invariance predicate*. It is read ‘When α is executed under condition ψ , the truth values of propositions p_1, \dots, p_m are invariant’.

Semantics

Standard modal logic structures (alias, possible worlds models) are tuples $\langle W, R, V \rangle$, where W is a (possibly infinite) set of states, R is a binary relation on W , and V is a valuation, assigning subsets of W to each atomic proposition. This is the standard Kripke-style semantics (Popkorn 1994; Hughes and Cresswell 1996, e.g.). SLAP structures are non-standard: Its semantics has a structure of the form $\langle W, R \rangle$, where W is a *finite* set of worlds such that each world assigns a truth value to each atomic proposition, and R is a binary relation on W . Moreover, SLAP is multi-modal in that there are multiple accessibility relations.

Intuitively, when talking about some world w , we mean 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*; the set of possible worlds may be the whole set of conceivable worlds.

Definition 2 A SLAP structure is a tuple $\mathcal{S} = \langle W, R \rangle$ s.t.

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 $\sum_{(w^-, w^+, pr) \in R_\alpha} pr = 0$;

Figure 1 is a pictorial representation of transitions and their probabilities for the action grab of the oil-can scenario.




Figure 1: A transition diagram for the grab action.

The eight circles represent the eight conceivable worlds with their valuations (\sim reads ‘not’).

R_α defines the transition probability $pr \in \mathbb{Q}_{[0,1]}$ between worlds w^+ and world 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^- .

Definition 3 (Truth Conditions) Let \mathcal{S} be a SLAP structure, with $\alpha \in \mathcal{A}$ and $q, pr \in \mathbb{Q}_{[0,1]}$. Let $p, p_1, \dots, p_m \in \mathcal{P}$ and let φ be any sentence in \mathcal{L}_{SLAP} . We say φ is satisfied at world w in structure \mathcal{S} (written $\mathcal{S}, w \models \varphi$) if and only if

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. $\mathcal{S}, w \models \neg\varphi \iff \mathcal{S}, w \not\models \varphi$;
4. $\mathcal{S}, w \models \varphi \wedge \varphi' \iff \mathcal{S}, w \models \varphi$ and $\mathcal{S}, w \models \varphi'$;
5. $\mathcal{S}, w \models [\alpha]_q\varphi \iff (\sum_{(w, w', pr) \in R_\alpha, \mathcal{S}, w' \models \varphi} pr) = q$;
6. $\mathcal{S}, w \models \Box\varphi \iff$ for all $w' \in W$, $\mathcal{S}, w' \models \varphi$;
7. $\mathcal{S}, w \models Inv(\alpha, \psi, \{p_1, \dots, p_m\}) \iff$
for $i \in \{1, \dots, m\}$, $\mathcal{S}, w \models \psi \wedge p_i \rightarrow [\alpha]p_i$ and
 $\mathcal{S}, w \models \psi \wedge \neg p_i \rightarrow [\alpha]\neg p_i$.

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

φ is *satisfiable* if $\mathcal{S}, w \models \varphi$ 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. We may thus write $w \models \varphi$ instead of $\mathcal{S}, w \models \varphi$ when φ is a propositional formula.

Let \mathcal{K} be a finite subset of \mathcal{L}_{SLAP} . We say that ψ is a *local semantic consequence* of \mathcal{K} (denoted $\mathcal{K} \models \psi$) if for all structures \mathcal{S} , and all $w \in W$ of \mathcal{S} , if $\mathcal{S}, w \models \bigwedge_{\theta \in \mathcal{K}} \theta$ then $\mathcal{S}, w \models \psi$. We shall also say that \mathcal{K} *entails* ψ whenever $\mathcal{K} \models \psi$. If $\{\theta\} \models \psi$ then we simply write $\theta \models \psi$.

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{P} .

Specifying Domains with SLAP

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 , that is, a specification of the world the agent finds itself in when it becomes active. (ii) Domain constraints or static laws SL , that is, facts and laws about the domain that do not change. (iii) Information about when actions are possible and impossible, the effects of actions and conditions for the effects—the dynamics of the environment or system. Refer to these as the *action description* (AD). How to write these axioms is the focus of this paper.

Let the union of all the axioms in SL and AD be denoted by the set BK —the agent’s *background knowledge*. IC is not part of the agent’s background knowledge.

Therefore, in SLAP we are interested in the validity of a formula with a particular form:

$$\models \bigwedge_{\phi \in BK} \Box \phi \rightarrow (IC \rightarrow \Phi), \quad (1)$$

where Φ is any sentence of interest in \mathcal{L}_{SLAP} . Equation (1) assumes that BK contains no sentence mentioning \Box .

From now on, the following abbreviations for constants in our scenario will be used: $\text{grab} := g$, $\text{drink} := d$, $\text{weigh} := w$, $\text{replace} := r$, $\text{full} := f$, $\text{drank} := da$ and $\text{holding} := h$.

Specifying Transition Models

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

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

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 4 *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*

$$\begin{aligned} \psi_1 &\rightarrow [\alpha]_{q_{11}} \varphi_{11} \wedge \dots \wedge [\alpha]_{q_{1n}} \varphi_{1n} \\ \psi_2 &\rightarrow [\alpha]_{q_{21}} \varphi_{21} \wedge \dots \wedge [\alpha]_{q_{2n}} \varphi_{2n} \\ &\vdots \\ \psi_j &\rightarrow [\alpha]_{q_{j1}} \varphi_{j1} \wedge \dots \wedge [\alpha]_{q_{jn}} \varphi_{jn}, \end{aligned}$$

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

We insist that no $q_{ik} = 0$, because the definition is of the specification of an action’s *effects*: suppose

$$\psi \rightarrow \dots \wedge [\alpha]_0 \varphi \wedge \dots$$

is an axiom of our background knowledge, then due to no φ -world being reachable via α under condition ψ , φ cannot

be an effect in this case. This axiom should thus not be an *effect* axiom.

Proper specifications of the probabilistic effects of actions g , d , r and w , respectively, are

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

$$\begin{aligned} f \wedge da \wedge h &\rightarrow [r](f \wedge da \wedge \neg h); \\ f \wedge \neg da \wedge h &\rightarrow [r](f \wedge \neg da \wedge \neg h); \\ \neg f \wedge da \wedge h &\rightarrow [r](\neg f \wedge da \wedge \neg h); \\ \neg f \wedge \neg da \wedge h &\rightarrow [r](\neg f \wedge \neg da \wedge \neg h). \end{aligned}$$

$$\begin{aligned} f \wedge da \wedge h &\rightarrow [w](f \wedge da \wedge h); \\ f \wedge \neg da \wedge h &\rightarrow [w](f \wedge \neg da \wedge h); \\ \neg f \wedge da \wedge h &\rightarrow [w](\neg f \wedge da \wedge h); \\ \neg f \wedge \neg da \wedge h &\rightarrow [w](\neg f \wedge \neg da \wedge h). \end{aligned}$$

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 ψ if there exists an effect axiom with condition ψ 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. Finally, one can only say what propositions do not change, under the conditions of the given axioms. However, a PES does not carry the information of whether the axioms are meant to cover *all* conditions. The rest of this paper is dedicated to dealing with these deficits.

If a knowledge engineer for some reason does not specify what an action α ’s effects are, given some condition ψ , but he/she wants to specify that the action is executable in ψ , then he/she can simply write $\psi \rightarrow [\alpha]_1 \top$. To express that α is inexecutable under condition ψ , the knowledge engineer can write $\psi \rightarrow [\alpha]_0 \top$.

Invariance

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 (1999) supply one approach to deal with the frame problem in a language able to express probabilistic transitions, but read the last section of the present paper.

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

$$\begin{aligned} h \wedge f &\rightarrow [r]f; & h \wedge \neg f &\rightarrow [r]\neg f; \\ h \wedge da &\rightarrow [r]da; & h \wedge \neg da &\rightarrow [r]\neg da. \end{aligned}$$

Here, h is the *condition* under which the frame axioms are applicable. An extreme example is for w (under condition h), which never affects propositions; six frame axioms are required. w is an epistemic action (without side-effects) because it is meant to affect only the agent's knowledge¹, not the environment.

In general, a positive frame axiom has the form

$$FrCond^+(\alpha, p) \wedge p \rightarrow [\alpha]p$$

and a negative frame axiom has the form

$$FrCond^-(\alpha, p) \wedge \neg p \rightarrow [\alpha]\neg p,$$

where $FrCond^+(\alpha, p)$ is a formula stating the conditions under which literal p remains positive and $FrCond^-(\alpha, p)$ is a formula stating the conditions under which literal $\neg p$ remains negative.

Instead of stating frame axioms directly, we shall use a slightly more concise *invariance predicate* by collecting all propositions invariant under the same conditions. To relate frame axioms and invariance predicates, note that the following two statements hold (\Rightarrow is read 'implies').

$$\begin{aligned} \mathcal{S}, w &\models Inv(\alpha, FrCond^+(\alpha, p) \wedge p, P) \text{ s.t. } p \in P \\ \Rightarrow \mathcal{S}, w &\models FrCond^+(\alpha, p) \wedge p \rightarrow [\alpha]p \end{aligned}$$

$$\begin{aligned} \mathcal{S}, w &\models Inv(\alpha, FrCond^-(\alpha, p) \wedge \neg p, P) \text{ s.t. } p \in P \\ \Rightarrow \mathcal{S}, w &\models FrCond^-(\alpha, p) \wedge \neg p \rightarrow [\alpha]\neg p. \end{aligned}$$

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 $Inv(\alpha, \psi, P), Inv(\alpha, \psi', P') \in INV$, $\psi \wedge \psi' \equiv \perp$. Furthermore, for every effect axiom $\psi \rightarrow \Phi$ for α , for all $Inv(\alpha, \psi', P) \in INV$, either $\psi \wedge \psi' \equiv \perp$ or $\psi \models \psi'$. These assumptions keep things organized.

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

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

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

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

¹In SLAP, actions like `weigh` are innocuous, but in logics which express sensing, epistemic actions become influential.

Note that $\bigwedge_{\beta \in PES_1} \beta \equiv \bigwedge_{\delta \in INV_1 \cup PES_2} \delta$, but $INV_1 \cup PES_2$ has approximately half as many formulae as PES_1 .

Furthermore, we propose that it is reasonable to deduce executability and inexecutability of actions by assuming the presence of the following (*in*)*executability axiom*.

$$\langle \alpha \rangle \top \leftrightarrow (\psi_1 \vee \dots \vee \psi_j) \vee \bigvee_{\psi \in Cond^{Inv}(\alpha)} \psi,$$

where ψ_1, \dots, ψ_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 (*in*)*executability axioms* for our example are

$$\begin{aligned} \langle g \rangle \top &\leftrightarrow \neg h; & \langle d \rangle \top &\leftrightarrow h \wedge (\neg f \vee \neg da); \\ \langle r \rangle \top &\leftrightarrow h; & \langle w \rangle \top &\leftrightarrow h. \end{aligned} \quad (2)$$

We shall collect all (*in*)*executability axioms* in the set $EXEC$. We shall refer to the set (2) in particular as $EXEC_1$.

From Underspecified to Fully Specified Transition Models

Definition 5 A transition model specification for action α is any set $B \subset \mathcal{L}_{SLAP}$, such that there exists a structure $\mathcal{S} = \langle C, R \rangle$, where $(\alpha, R_\alpha) \in R$ such that $\mathcal{S} \models \bigwedge_{\beta \in B} \beta$ 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 $R'_\alpha \neq R_\alpha$.

A PES is, in general, not a transition model specification: Let $\mathcal{P} = \{p_1\}$, $\alpha_1 \in \mathcal{A}$ 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 $\mathcal{S}' = \langle C, R' \rangle$, $(w_2, w_2, 0.4) \in R'_\alpha$ and $(\alpha, R') \in R'$ and assume $\mathcal{S}'' \models p_1 \rightarrow [\alpha_1]p_1$, where $\mathcal{S}'' = \langle C, R'' \rangle$, $(w_2, w_2, 0.5) \in R''_\alpha$ and $(\alpha, R'') \in R''$. 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 transition model specification 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 proposition to change (see, e.g., 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 $BK^{ocs} := INV_1 \cup PES_2 \cup EXEC_1$. Note that $BK^{ocs} \not\models f \wedge \neg h \rightarrow [g]_q(f \wedge da \wedge \neg h) \vee [g]_{q'}(\neg f \wedge da \wedge \neg h)$ for any q and q' . One could assume, due to lack of knowledge, that the truth value of f does not change, that is

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

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

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

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.

Definition 6 Given effect axiom $\psi \rightarrow [\alpha]_{q_1} \varphi_1 \wedge [\alpha]_{q_2} \varphi_2 \wedge \dots \wedge [\alpha]_{q_n} \varphi_n$ for α of a PES, proposition p is effectively underspecified in effect $\varphi \in \{\varphi_1, \varphi_2, \dots, \varphi_n\}$ under condition ψ if and only if $[\alpha]_q \varphi \not\models [\alpha]_q (\varphi \wedge p)$ and $[\alpha]_q \varphi \not\models [\alpha]_q (\varphi \wedge \neg p)$ and if there exists $Inv(\alpha, \psi', P) \in INV$ such that $\psi \models \psi'$, then $p \notin P$.

Definition 7 Given invariance predicate $Inv(\alpha, \psi, P) \in INV$, proposition p is completely invariantly underspecified under condition ψ if and only if $p \notin P$, and there exists no condition ψ' of an effect axiom for α of the PES, such that $\psi' \models \psi$.

Definition 8 Given invariance predicate $Inv(\alpha, \psi, P) \in INV$, proposition p is partially invariantly underspecified under condition ψ if and only if $p \notin P$, and there exists condition ψ' of an effect axiom for α of the PES, such that $\psi' \models \psi$, but $\psi' \not\models \psi$.

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 proposition is *underspecified*, we mean it in the sense of one or more of Definitions 6, 7 or 8. We propose two alternative approaches, investigated formally in the subsections below. When a proposition 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.

Always Assuming Invariance

Let ψ^{prtl} be $\psi \wedge \neg \psi'$, where ψ' is the condition of an effect axiom for α of the PES, and $Inv(\alpha, \psi, P) \in INV$, such that $\psi' \models \psi$, but $\psi' \not\models \psi$.

This approach is: for every $\alpha \in \mathcal{A}$ and $p \in \mathcal{P}$, if p is effectively underspecified in effect φ under condition ψ for α , assume the presence of *invariance formulae*

$$\psi \wedge p \rightarrow [\alpha]_q (\varphi \wedge p) \text{ and } \psi \wedge \neg p \rightarrow [\alpha]_q (\varphi \wedge \neg p),$$

else if p is completely invariantly underspecified under condition ψ , add p to P of $Inv(\alpha, \psi, P) \in INV$, else if p is partially invariantly underspecified under condition ψ , assume the presence of *invariance formulae*

$$\psi^{prtl} \wedge p \rightarrow [\alpha]_q p \text{ and } \psi^{prtl} \wedge \neg p \rightarrow [\alpha]_q \neg p.$$

Given INV_1 and PES_2 , the invariance formula and new invariance predicate assumed present are

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

$$Inv(da, \neg f \wedge \neg da \wedge h, \{f, da, h\}).$$

Given PES_1 , the invariance formulae assumed present are

$$f \wedge da \wedge \neg h \rightarrow [g]_{0.3} (da \wedge \neg h \wedge f);$$

$$f \wedge \neg da \wedge \neg h \rightarrow [g]_{0.3} (\neg da \wedge \neg h \wedge f);$$

$$\neg f \wedge \neg da \wedge h \rightarrow [d] (\neg f \wedge \neg da \wedge h).$$

Always Assuming Uniform Distribution

Let $U^e(\alpha, \psi, \varphi) = \{p \in \mathcal{P} \mid p \text{ is effectively underspecified for } \alpha \text{ in effect } \varphi \text{ under condition } \psi\}$, $U^c(\alpha, \psi) = \{p \in \mathcal{P} \mid p \text{ is completely invariantly underspecified for } \alpha \text{ under condition } \psi\}$ and $U^p(\alpha, \psi) = \{p \in \mathcal{P} \mid p \text{ is partially invariantly underspecified for } \alpha \text{ under condition } \psi\}$.

Then this approach is: for every action α , for every transition $[\alpha]_q \varphi$ of every effect axiom with condition ψ , assume the presence of *equiprob formula*

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

where $\{\gamma_1, \dots, \gamma_m\}$ are the $m = 2^{|U^e(\alpha, \psi, \varphi)|}$ permutations of conjunctions of literals, given all the propositions in $U^e(\alpha, \psi, \varphi)$ and $q_1 = \dots = q_m = q/m$. For instance, if $U^e(\alpha, \psi, \varphi) = \{p_2, p_4\}$ then the literal conjunction permutations are $\{p_2 \wedge p_4, p_2 \wedge \neg p_4, \neg p_2 \wedge p_4, \neg p_2 \wedge \neg p_4\}$.

Else, for every action α , for every invariance predicate with condition ψ , assume the presence of *equiprob formula*

$$\psi \rightarrow [\alpha]_{q_1} \gamma_1 \wedge \dots \wedge [\alpha]_{q_n} \gamma_n,$$

where $\{\gamma_1, \dots, \gamma_n\}$ are the $n = 2^{|U^c(\alpha, \psi)|}$ permutations of conjunctions of literals, given all the propositions in $U^c(\alpha, \psi)$ and $q_1 = \dots = q_n = q/n$, and similarly for $U^p(\alpha, \psi)$, but with condition ψ^{prtl} for the equiprob formula.

Given INV_1 and PES_2 , the following equiprob formulae are assumed present.²

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

$$\neg f \wedge \neg da \wedge h \rightarrow [d]_{0.5} da \wedge [d]_{0.5} \neg da.$$

Given PES_1 , the following equiprob formulae are assumed present.

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

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

$$\neg f \wedge \neg da \wedge h \rightarrow [d]_{0.5} (\neg f \wedge da \wedge h) \wedge [d]_{0.5} (\neg f \wedge \neg da \wedge h).$$

The Two Approaches are Full Specifications

Definition 9 A verbose effects specification (*VES*) is a PES where all effect axiom conditions (the ψ left of the \rightarrow) and effects (the φ right of the \rightarrow) are *definitive formulae*.

Lemma 1 Let E^V be $\langle \alpha \rangle \top \leftrightarrow (\psi_1 \vee \dots \vee \psi_j)$, where ψ_1, \dots, ψ_j are the j conditions of the j effect axioms in a VES V for α . Then $E^V \wedge \bigwedge_{\beta \in V} \beta$ is a transition model specification.

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 $\langle \alpha, R_\alpha \rangle \in R$ and $\langle C, R \rangle \models E^V \wedge \bigwedge_{\beta \in V} \beta$.

²Note that $\{\neg f \wedge \neg da \wedge h \rightarrow [d] (\neg f \wedge h), \neg f \wedge \neg da \wedge h \rightarrow [d]_{0.5} da \wedge [d]_{0.5} \neg da\} \models \neg f \wedge \neg da \wedge h \rightarrow [d]_{0.5} (\neg f \wedge da \wedge h) \wedge [d]_{0.5} (\neg f \wedge \neg da \wedge h)$.

For the sake of reference, let

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

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 $w^- \not\models (\psi_1 \vee \dots \vee \psi_j)$, then $(w^-, w^+, 0) \in R_\alpha$. Else if $w^- \models \psi$: if $w^+ \models \varphi_k$ then $(w^-, w^+, q_k) \in R_\alpha$, else if $w^+ \not\models \varphi_1 \vee \dots \vee \varphi_n$ 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 $\varphi_{ik'}$, $\varphi_{ik} \wedge \varphi_{ik'} \equiv \perp$, 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 ψ_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}, \dots, q_{in} of any axiom i must sum to 1. Else, for all $w^- \in C$ such that $w^- \models \neg(\psi_1 \vee \dots \vee \psi_j)$, $(w^-, w^+, 0) \in R_\alpha$ for all $w^+ \in C$, due to the PES constraint that for any pair of conditions ψ_i and $\psi_{i'}$, $\psi_i \wedge \psi_{i'} \equiv \perp$. 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 $\langle C, R \rangle \models E^V$. 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 $\langle C, R' \rangle \models E^V \wedge \bigwedge_{\beta \in V} \beta$. 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 $w^- \not\models (\psi_1 \vee \dots \vee \psi_j)$, then $(w^-, w^+, q') \in R'_\alpha$, where $q' > 0$. But then $\langle C, R' \rangle \not\models E^V$. And if $w^- \models \psi$ and $w^+ \models \varphi_k$, then $q_k \neq q'$ and $\langle C, R' \rangle \not\models \psi \rightarrow [\alpha]_{q_k} \varphi_k$, 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 $w^- \models \psi$ and $w^+ \not\models \varphi_1 \vee \dots \vee \varphi_n$ 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}, \dots, q_{in} of any axiom i must sum to 1, $\sum_{(w^-, w^+, pr) \in R_\alpha} pr > 1$. ■

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

Proof:

(Abridged) Let \mathcal{S} be an arbitrary SLAP structure and w a world in it. Suppose $\mathcal{S}, w \models (\psi \rightarrow [\alpha]_q \varphi) \wedge (\psi' \rightarrow [\alpha]_{q'} \varphi') \wedge (\psi' \rightarrow \psi)$. Assume $\mathcal{S}, w \models \psi'$. Then $\mathcal{S}, w \models [\alpha]_q \varphi \wedge [\alpha]_{q'} \varphi'$.

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

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

Hence, $\mathcal{S}, w \models$ iff $\sum_{(w, w', pr) \in R_\alpha, \mathcal{S}, w' \models \varphi \wedge \varphi'} pr = q$ iff $\mathcal{S}, w \models [\alpha]_q (\varphi \wedge \varphi')$. ■

Theorem 1 For both approaches, given a PES Pes for α , a set of invariance predicates INV for α , a set of (in)executability axioms E for α derived from Pes and

INV , a set of invariance formulae IF for α and a set of equiprob formulae EF for α , their union is a transition model specification.

Proof:

(Abridged) Suppose V is a VES for α and E^V is an (in)executability axiom derived from V as in Lemma 1. If we can show that V and E^V exist such that $E \wedge \bigwedge_{\beta \in Pes \cup INV \cup IF \cup EF} \beta \equiv E^V \wedge \bigwedge_{\delta \in V} \delta$, then by Lemma 1, we have proved the theorem. Hence, we show how to convert $Pes \cup INV \cup IF \cup EF$ into a semantically equivalent VES V and we prove that $E \equiv E^V$.

With several rewrite rules, we show how to ‘enlarge’ Pes into a VES using the information in INV , IF and EF . The following rule is applied to both default assumption approaches. • For every $Inv(\alpha, \psi, P) \in INV$, add $\psi \wedge p \rightarrow [\alpha]_p$ and $\psi \wedge \neg p \rightarrow [\alpha]_{\neg p}$ to Pes , for every $p \in P$. Next, considering the two approaches separately.

Suppose *Always Assuming Invariance* is used.

- Note that EF is empty.
- Let $L^{IF}(\alpha, \psi, \varphi) = \{\ell \mid \psi \wedge \ell \rightarrow [\alpha]_q (\varphi \wedge \ell) \in IF\}$. For every $\psi \rightarrow [\alpha]_{q_1} \varphi_1 \wedge [\alpha]_{q_2} \varphi_2 \wedge \dots \wedge [\alpha]_{q_n} \varphi_n$ in Pes , replace φ_k by $\varphi_k \wedge \bigwedge_{\ell \in L^{IF}(\alpha, \psi, \varphi_k)} \ell$.
- Replace $\psi \rightarrow \Phi$ by the members of the set $\{\delta \rightarrow \Phi \mid \delta \models \psi, \delta \in Def\}$.
- By Proposition 1, for every $\delta \rightarrow [\alpha]_{q_1} \varphi_1 \wedge [\alpha]_{q_2} \varphi_2 \wedge \dots \wedge [\alpha]_{q_n} \varphi_n$ in Pes , replace φ_k by $\varphi_k \wedge \bigwedge_{\ell \in \{\ell \mid \delta \rightarrow [\alpha]_{\ell} \in Pes\}} \ell$.
- Remove all formulae of the form $\psi \rightarrow [\alpha]_{\ell}$ from Pes (we assume that the vocabulary has > 1 proposition).
- Remove all but one semantically equivalent formulae from Pes . It should be easy to recognize which formulae are equivalent, given the form they are now in.

Suppose *Always Assuming Uniform Distribution* is used.

- Note that IF is empty.
- Add all members of EF to Pes .
- Replace $\psi \rightarrow \Phi$ by the members of the set $\{\delta \rightarrow \Phi \mid \delta \models \psi, \delta \in Def\}$.
- By Proposition 1, for every $\delta \rightarrow [\alpha]_{q_1} \varphi_1 \wedge [\alpha]_{q_2} \varphi_2 \wedge \dots \wedge [\alpha]_{q_n} \varphi_n$ in Pes , replace φ_k by $\varphi_k \wedge \bigwedge_{\ell \in \{\ell \mid \delta \rightarrow [\alpha]_{\ell} \in Pes\}} \ell$.
- For every $\delta \rightarrow [\alpha]_{q_1} \varphi_1 \wedge [\alpha]_{q_2} \varphi_2 \wedge \dots \wedge [\alpha]_{q_n} \varphi_n$ in Pes , for every other $\delta' \rightarrow [\alpha]_{q'_1} \varphi'_1 \wedge [\alpha]_{q'_2} \varphi'_2 \wedge \dots \wedge [\alpha]_{q'_n} \varphi'_n$ in Pes , replace $[\alpha]_{q_k} \varphi_k$ by $\bigwedge_{\varphi'_k \models \varphi_k} [\alpha]_{q'_k} \varphi'_k$.
- Remove all formulae of the form $\psi \rightarrow [\alpha]_{\ell}$ (we assume that the vocabulary has > 1 proposition) and of the form $\psi \rightarrow [\alpha]_{q_1} \varphi_1 \wedge [\alpha]_{q_2} \varphi_2 \wedge \dots \wedge [\alpha]_{q_n} \varphi_n$ s.t. $\sum_{i=1}^n q_i < 1$ from Pes .
- Remove all but one semantically equivalent formulae from Pes . It should be easy to recognize which formulae are equivalent, given the form they are now in.

By the nature of INV , IF and EF , every effect of every effect axiom is now a definitive formula.

Observe that E depends only on the axiom conditions of the original Pes , which has essentially the same axiom conditions as those of V (given our assumption that all effect axiom conditions in Pes are definitive), and E^V depends only on the axiom conditions of V . Hence $E \equiv E^V$. ■

Discussion and Related Work

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 managing the ‘full’ specification of probabilistic transition models 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, he/she must take a stance as to what the default transition behavior should be: invariance of the truth values of propositions 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 proposition should change when some action is executed, it could be assumed that the proposition’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 nonmonotonic reasoning (Brewka 2012).

The ‘uniform distribution’ assumption on the other hand is a kind of ‘maximum entropy’ approach. Wang and Schmolze (2005) have a very similar approach to ours to achieve compact representations in POMDP planning. Some researchers (see, e.g., (Grove, Halpern, and Koller 1994) and the work of 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.* (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 (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 nondeterministic action by ‘decomposing’ it into a set of deterministic actions, each leading to one of the effects of the nondeterministic action. We opted to specify stochastic (nondeterminis-

tic) actions ‘directly’ and not to decompose them. It is our opinion that our ‘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. Without a notion of equality between actions, one cannot write SSAs in SLAP. Nevertheless, even with such a notion, 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

- Bacchus, F.; Halpern, J. Y.; and Levesque, H. J. 1999. Reasoning about noisy sensors and effectors in the situation calculus. *Artificial Intelligence* 111(1–2):171–208.
- Brewka, G. 2012. *Nonmonotonic Reasoning: Logical Foundations of Commonsense (Cambridge Tracts in Theor. Comp. Sci.)*. Cambridge University Press, reissue edition.
- Fierens, D.; Blockeel, H.; Bruynooghe, M.; and Ramon, J. 2005. Logical bayesian networks and their relation to other probabilistic logical models. In Kramer, S., and Pfahringer, B., eds., *ILP 2005*, volume 3625 of *LNAI*, 121–135. Springer Verlag.
- Grove, A.; Halpern, J.; and Koller, D. 1994. Random worlds and maximum entropy. *Journal of Artificial Intelligence Research (JAIR)* 2:33–88.
- Hughes, G., and Cresswell, M. 1996. *A New Introduction to Modal Logic*. New York, NY: Routledge.
- Jaeger, M. 1994. A logic for default reasoning about probabilities. In *Proc. of 10th Intl. Conf. on Uncertainty in Artif. Intell.*, UAI’94, 352–359. Morgan Kaufmann Publishers Inc.
- Jaynes, E. 1978. Where do we stand on maximum entropy? In *The Maximum Entropy Formalism*. MIT Press. 15–118.
- Kern-Isberner, G. 2001. Conditionals in nonmonotonic reasoning and belief revision. In *Lecture Notes in Computer Science (LNCS)*, number 2087. Springer Verlag.
- McCarthy, J., and Hayes, P. 1969. Some philosophical problems from the standpoint of artificial intelligence. *Machine Intelligence* 4:463–502.
- Popkorn, S. 1994. *First Steps in Modal Logic*. Cambridge University Press.
- Reiter, R. 1991. The frame problem in the situation calculus: A simple solution (sometimes) and a completeness result for goal regression. In Lifschitz, V., ed., *Artificial intelligence and mathematical theory of computation: papers in honor of John McCarthy*. San Diego, CA, USA: Academic Press Professional, Inc. 359–380.
- Wang, C., and Schmolze, J. 2005. Planning with POMDPs using a compact, logic-based representation. In *Proc. of 17th IEEE Intl. Conf. on Tools with Artif. Intell. (ICTAI’05)*, 523–530. Los Alamitos, CA, USA: IEEE Computer Society.