Journal of Applied Logic 12 (2014) 128–150

Contents lists available at ScienceDirect

Journal of Applied Logic

www.elsevier.com/locate/jal

SLAP: Specification logic of actions with probability

Gavin Rens [a,*], Thomas Meyer [a], Gerhard Lakemeyer [b]

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

ARTICLE INFO

Article history:
Received 7 February 2013
Accepted 25 September 2013
Available online 23 October 2013

Keywords:
Probabilistic actions
Modal logic
Tableau method
Systems of linear inequalities

ABSTRACT

A logic for specifying probabilistic transition systems is presented. Our perspective is that of agents performing actions. A procedure for deciding whether sentences in this logic are valid is provided. One of the main contributions of the paper is the formulation of the decision procedure: a tableau system which appeals to solving systems of linear equations. The tableau rules eliminate propositional connectives, then, for all open branches of the tableau tree, systems of linear equations are generated and checked for feasibility. Proofs of soundness, completeness and termination of the decision procedure are provided.

1. Introduction

In this article, we present a logic for specifying agents' stochastic action models, or more generally, for specifying probabilistic transition systems—the Specification Logic of Actions with Probability (SLAP). Our logic takes the possible worlds semantics of modal logic and draws inspiration from Markov decision process (MDP) theory to deal with probabilities.

Modal logic [20,10,4,5] is considered to be well suited to reasoning about beliefs and changing situations and MDP theory [3,19,31] has proven to be a good general framework for formalizing dynamic stochastic systems.

Next we introduce a scenario to illustrate concepts throughout the article. 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 $\mathcal{A} = \{\mathtt{grab}, \mathtt{drink}, \mathtt{weigh}, \mathtt{replace}\}$ with expected

* Corresponding author.
E-mail addresses: grens@csir.co.za (G. Rens), tmeyer@csir.co.za (T. Meyer), gerhard@cs.rwth-aachen.de (G. Lakemeyer).

1570-8683/$ – see front matter
http://dx.doi.org/10.1016/j.jal.2013.09.001

meanings. The robot experiences its world (domain) through three Boolean features: $\mathcal{P} = \{\mathtt{full}, \mathtt{drank}, \mathtt{holding}\}$ meaning respectively 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 queries:

- If the oil-can is full, I have not drunk the contents and I am holding the can, is there a 0.15 probability that after 'drinking' the contents, the oil-can is still full, I have still not drunk the oil and I am still holding the can? That is, does $(\mathtt{full} \land \neg\mathtt{drank} \land \mathtt{holding}) \to [\mathtt{drink}]_{0.15}(\mathtt{full} \land \neg\mathtt{drank} \land \mathtt{holding})$ follow from BK?
- If the oil-can is empty and I'm not holding it, is there a 0.9 probability that I'll be holding it after grabbing it, and a 0.1 probability that I'll have missed it? That is, does $(\neg\mathtt{full} \land \neg\mathtt{holding}) \to ([\mathtt{grab}]_{0.9}(\neg\mathtt{full} \land \mathtt{holding}) \land [\mathtt{grab}]_{0.1}(\neg\mathtt{full} \land \neg\mathtt{holding}))$ follow from BK?

In another paper [34], we propose how to write sentences in the language of SLAP to capture a model of probabilistic transitions due to the execution of actions of some agent. And in that paper, we suggest which assumptions can and perhaps should be made about such specifications to make them more parsimonious. 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 that paper [34].

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.

In SLAP we are interested in whether $IC \to \Phi$ follows from $\bigwedge_{\phi \in BK} \Box\phi$, where Φ is any 'legal' sentence of interest and $\Box$ marks sentences as laws of the domain, that is, sentences which must be true in every possible world.

In Section 2.2, we define *entailment*, which depends on the notion of *validity*. We must defer a discussion of the use of SLAP to Section 2.3, after the syntax and semantics have been presented. The focus of this article, though, is on a decision procedure for entailment of SLAP sentences from sets of SLAP sentences, and on the computational property of the decision procedure.

Section 2 defines SLAP. Section 3 provides a decision procedure for determining entailment of sentences in SLAP. In Section 4, we prove that the procedure is sound, complete and that it terminates, that is, we show that SLAP is decidable with respect to entailment. Sections 5 and 6 cover some related work, and respectively, summarizes what has been achieved in this article, and discusses future work.

2. Specification logic of actions with probability

First we present the syntax of SLAP, then we state its semantics.

2.1. Syntax

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* $\mathbb{Q}$.