A Logic for Reasoning about Actions and Explicit Observations

Gavin Rens[1,2], Ivan Varzinczak[1], Thomas Meyer[1,2], and Alexander Ferrein[3]

[1] KRR, CSIR Meraka, Pretoria, South Africa
[2] School of Computer Science, University of KwaZulu-Natal, South Africa
{grens,ivarzinczak,tmeyer}@meraka.org.za
[3] Robotics and Agents Research Lab, University of Cape Town, South Africa
alexander.ferrein@uct.ac.za

Abstract. We propose a formalism for reasoning about actions based on multi-modal logic which allows for expressing observations as first-class objects. We introduce a new modal operator, namely $[o \mid \alpha]$, which allows us to capture the notion of perceiving an observation given that an action has taken place. Formulae of the type $[o \mid \alpha]\varphi$ mean 'after perceiving observation o, given α was performed, necessarily φ'. In this paper, we focus on the challenges concerning sensing with explicit observations, and acting with nondeterministic effects. We present the syntax and semantics, and a correct and decidable tableau calculus for the logic.

1 Introduction and Motivation

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 arm. 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 $\mathfrak{A} = \{grab, drink, weigh, replace\}$ with expected meanings. The robot can perceive observations only from the set $\Omega = \{obsNil, obsHeavy, obsMedium, obsLight\}$. Intuitively, when the robot performs a *weigh* action, it will perceive either *obsHeavy*, *obsMedium* or *obsLight*; for other actions, it will 'perceive' *obsNil*, no perception. The robot experiences its world (domain) via three Boolean features: $\mathfrak{P} = \{full, drank, holding\}$ meaning respectively that the the oil-can is full, that the robot has drunk the oil and that it is currently holding something in its gripper. This formalization seems more intuitive than lumping all observations in with propositions, for instance, by making $\mathfrak{P} = \{full, drank, holding, obsnil, heavy, medium, light\}$.

It is the norm in dynamic logics (and some other agent oriented logics) to deal with observations as elements of knowledge, as propositions; and perception is normally coded as action, that is, observations-as-propositions evaluate to 'true'

J. Li (Ed.): AI 2010, LNAI 6464, pp. 395–404, 2010.

or 'false' depending on some action(s). However, the approach of interpreting observations as mere propositions may be counterintuitive to some people, because knowledge may be seen as something different from events (observations) that *generate* or *modify* knowledge.

Remark 1. If an intelligent agent is regarded as a *system*, then there are inputs to the system that affect it, and outputs from the system that affect the environment. The inputs are *observations* and the outputs are *actions*. If one assumes that the system state is represented by a knowledge base of *propositions*, then from the systems view, it is clear that observations and propositions are different in nature.

Therefore, the ability to distinguish between observations and propositions allows for a more precise specification of a given domain, as we shall see in the sequel. It turns out that the notion of observations as explicit syntactic and semantic objects of a logic is not completely new. For example, Van Benthem, Gerbrandy and Kooi [15] do so (For more details, the reader is invited to see Section 5 on related work.)

Although there are several formalisms in the literature on reasoning about and specifying agents and their actions, we found them lacking when it comes to treating observations as objects on a par with actions, while retaining important computational properties. Existing first-order based approaches are in general undecidable or have too complicated semantics. For these reasons, we prefer to anchor our framework on a version of dynamic logic and strive for an extension of it by allowing for observations as explicit entities.

The rest of this paper is organized as follows. We give the syntax and semantics of our logic in Section 2. In Section 3, we show how to correctly specify agent domains with our logic. Our tableau method, with correctness and decidability results, is given in Section 4. Section 5 covers related work and Section 6 concludes the paper.

2 A Logic for Actions and Observations

The logic we present here allows for expressing observations explicitly, distinct from propositions. It is called the Logic for Actions and Observations (LAO). LAO is a non-standard modal logic with quantification and equality over the actions and observations. It will be able to accommodate formal descriptions of nondeterminism in the actions and of uncertainty in the observations. Given a formalization $\mathcal{K}$ of our scenario, the robot may have the following queries:

- Is it possible that after grabbing the oil-can, I will not be holding it? That is, does $\langle grab \rangle \neg holding$ follow from $\mathcal{K}$?
- If I weigh the oil-can and perceive that it is heavy, is it necessary that I have drunk the oil? That is, does $[obsHeavy \mid weigh]\, drank$ follow from $\mathcal{K}$?

LAO is based on $\mathcal{LAP}$ (the Logic for Actions and Plans [1]), but with one major difference: the addition of *observations*. That is, LAO refers to a set of observations that are explicitly identified by a knowledge engineer or agent-system designer (cf. Remark 1). A minor, yet important difference is the addition of action and observation variables, quantification and equality.

2.1 Syntax

We work in a propositional language. It contains three sorts: (1) a finite set of *fluents* (alias *propositional atoms*) $\mathfrak{P} = \{p_1, \ldots, p_n\}$, (2) a finite set of names of atomic *actions* $\mathfrak{A} = \{\alpha_1, \ldots, \alpha_n\}$ and a countable set of *action variables* $V_\mathfrak{A} = \{v_1^\alpha, v_2^\alpha, \ldots\}$, and (3) a finite set of names of atomic *observations* $\Omega = \{\varsigma_1, \ldots, \varsigma_n\}$ and a countable set of *observation variables* $V_\Omega = \{v_1^\varsigma, v_2^\varsigma, \ldots\}$. We shall refer to elements of $\mathfrak{A} \cup \Omega$ as *constants* and elements of $V_\mathfrak{A} \cup V_\Omega$ as *variables.* A *literal* ℓ is a fluent or its negation.

We are going to work in a multi-modal setting, in which we have a modal operator $[\alpha]$, one for each element in $\mathfrak{A}$; and a modal operator $[\varsigma|\alpha]$, one for each pair (α, ς) in $\mathfrak{A} \times \Omega$.

Definition 1. *Let $\alpha, \alpha' \in (\mathfrak{A} \cup V_\mathfrak{A})$, $\varsigma, \varsigma' \in (\Omega \cup V_\Omega)$, $v \in (V_\mathfrak{A} \cup V_\Omega)$ and $p \in \mathfrak{P}$. The language of LAO, denoted $\mathcal{L}_{LAO}$, is the least set of those φ that contain no free variables:*

$$\varphi ::= p \mid \top \mid \neg\varphi \mid \varphi \wedge \varphi \mid \alpha = \alpha' \mid \varsigma = \varsigma' \mid [\alpha]\varphi \mid [\varsigma \mid \alpha]\varphi \mid (\forall v)\varphi.$$

For example, $[v^\varsigma|\alpha]$ is not in $\mathcal{L}_{LAO}$, but $(\forall v^\varsigma)[v^\varsigma|\alpha]$ is.

As usual, we treat $\bot, \vee, \rightarrow, \leftrightarrow, \neq$ and $\exists$ as abbreviations. The sentence $[\varsigma \mid \alpha]\varphi$ is read 'φ must hold after ς is observed, given α is executed'. For instance, $[\mathit{obsLight} \mid \mathit{weigh}]\neg\mathit{full}$ means 'After perceiving that the oil-can is light, given a weighing action, the can is necessarily not full'. $[\alpha]\varphi$ is read 'φ must hold (after any/every observation) given α is executed'. For instance, $[\mathit{replace}]\neg\mathit{holding}$ means 'After replacing the oil-can, it is definitely not being held (regardless of observations)'. $\langle\alpha\rangle\varphi$ and $\langle\varsigma \mid \alpha\rangle\varphi$ abbreviate $\neg[\alpha]\neg\varphi$ and $\neg[\varsigma \mid \alpha]\neg\varphi$ respectively. One conventional reading for $\langle\alpha\rangle\varphi$ is 'It is possible that φ holds after α is performed'. The reading of $\langle\varsigma \mid \alpha\rangle\varphi$ is 'It is possible that φ holds after ς is perceived, given α is performed'.

We say that a formula is *static* if it mentions no actions.

We write $\varphi|_c^v$ to mean the formula φ with all variables v appearing in it replaced by constant c of the right sort (action or observation).

2.2 Semantics

Our semantics follows that of multi-modal logic K [11]. However, structures (alias, possible worlds models) are non-standard. Intuitively, when talking about some world w, we mean a set of features (*fluents*) that the agent understands and that describes a state of affairs in the world or that describes a possible, alternative world. Let $w : \mathfrak{P} \longrightarrow \{0, 1\}$ be a total function that assigns a truth value to each fluent. Let S be the set of all possible functions w. We call S the *conceivable worlds.*

Definition 2. *A LAO structure is a tuple $\mathcal{S} = \langle W, R, O, N, Q\rangle$ such that*

1. *$W \subseteq S$ is a non-empty (finite) set of* possible worlds*;*
2. *R is a mapping that provides an accessibility relation $R_\alpha : W \longrightarrow W$ for each action $\alpha \in \mathfrak{A}$;*
3. *O is a non-empty finite set of observations;*
4. *$N : \Omega \longrightarrow O$ is a total bijection that associates to each name in Ω, a unique observation in O;*
5. *Q is a mapping that provides a perceivability relation $Q_\alpha : O \longrightarrow W$ for each action $\alpha \in \mathfrak{A}$;*
6. *For all w, w', α, if $(w, w') \in R_\alpha$ then there is an o s.t. $(o, w') \in Q_\alpha$, for $w, w' \in W$, $\alpha \in \mathfrak{A}$ and $o \in O$.*

R_α defines which worlds w^+ are accessible via action α performed in world w^- and Q_α defines which observations o are perceivable in worlds w^+ accessible via action α. For $\varsigma \in \Omega$, $N(\varsigma) = o \in O$. Because N is a total bijection, it follows that $|O| = |\Omega|$.

Item 6 of Definition 2 implies that actions and observations always appear in pairs, even if implicitly. For example, if action *open-eyes* is performed, several signals are possible, depending on the situation, like *wall-3-meters-ahead* or *overcast-sky*. If the agent performs an action like *step-once-forward*, there is only one observation possible, viz. *null*, the 'dummy' observation. Unlike the eye, the leg (or wheel) is not a sensory organ. When our agent activates a device (the agent acts) and the device receives no input signal, it interprets this state of affairs as the *null* observation (the *null* observation will be denoted by the special named constant, *obsNil*). For every action an agent performs, the agent perceives exactly one observation. This is the approach of POMDPs that we rely on in the present work [10].

Definition 3 (Truth Conditions). *Let $\mathcal{S}$ be a LAO structure, with $\alpha, \alpha' \in \mathfrak{A}$, $v^\alpha \in V_{\mathfrak{A}}$, $\varsigma, \varsigma' \in \Omega$, $v^\varsigma \in V_\Omega$ and $p \in \mathfrak{P}$. And let φ be any sentence in $\mathcal{L}_{LAO}$. We say φ is* satisfied *at world w in structure $\mathcal{S}$ (written $\mathcal{S}, w \models \varphi$):*

1. $\mathcal{S}, w \models p$ iff $w(p) = 1$;
2. $\mathcal{S}, w \models \top$ for any $w \in W$;
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. $\mathcal{S}, w \models \alpha = \alpha'$ iff $\alpha, \alpha' \in \mathfrak{A}$ are the same element;
6. $\mathcal{S}, w \models \varsigma = \varsigma'$ iff $\varsigma, \varsigma' \in \Omega$ are the same element;
7. $\mathcal{S}, w \models [\alpha]\varphi$ iff for all w' and o, if $(w, w') \in R_\alpha$ and $(o, w') \in Q_\alpha$ then $\mathcal{S}, w' \models \varphi$;
8. $\mathcal{S}, w \models [\varsigma \mid \alpha]\varphi$ iff for all w', if $(w, w') \in R_\alpha$ and $(N(\varsigma), w') \in Q_\alpha$ then $\mathcal{S}, w' \models \varphi$;
9. $\mathcal{S}, w \models (\forall v^\alpha)\varphi$ iff $\mathcal{S}, w \models \varphi|^{v^\alpha}_{\alpha}$ for all $\alpha \in \mathfrak{A}$;
10. $\mathcal{S}, w \models (\forall v^\varsigma)\varphi$ iff $\mathcal{S}, w \models \varphi|^{v^\varsigma}_{\varsigma}$ for all $\varsigma \in \Omega$.

A formula φ is true (valid) in a LAO structure (denoted $\mathcal{S} \models \varphi$) if $\mathcal{S}, w \models \varphi$ for every $w \in W$. φ is LAO-valid (denoted $\models_{LAO} \varphi$) if φ is true in every structure $\mathcal{S}$. φ is *satisfiable* if $\mathcal{S}, w \models \varphi$ for some $\mathcal{S}$ and $w \in W$. We define *global logical entailment* (denoted $\psi \models_G \varphi$) as follows: for all $\mathcal{S}$, if $\mathcal{S} \models \psi$, then $\mathcal{S} \models \varphi$.

The motivation behind the definition of $\mathcal{S}, w \models [o \mid \alpha]\varphi$ is as follows. Just as φ needs not hold in worlds w' if $(w, w') \notin R_\alpha$, worlds w' are not considered if $(o, w') \notin Q_\alpha$. In other words, whether or not a world w' is reachable (via R_α), if the agent perceived o and the agent knows that o is not perceivable in w', then the agent knows it is not in w'. Then what is true or false in w' has no influence on the verity of $\mathcal{S}, w \models [o \mid \alpha]\varphi$. But in every world w' reachable from w *and* in which o is perceivable, φ must be true. While actions can add worlds that an agent believes possible, thus increasing uncertainty, observations eliminate reachable worlds from consideration, thus increasing certainty.

Proposition 1. $\models_{LAO} (\forall v^\alpha)\langle v^\alpha\rangle\varphi \rightarrow (\exists v^\varsigma)\langle v^\varsigma \mid v^\alpha\rangle\varphi$.

This means that for any structure $\mathcal{S}$ and world w, for any action α, if world w' can be reached from w via α, then there exists an observation perceivable in w'. Proposition 1 follows from item 6 of Definition 2.

Due to the nature of the 'observation naming' function N, in the rest of this paper, in our intuitive explanations, we let o mean o or ς (such that $N(\varsigma) = o$) depending on the context, and similarly we let ς mean ς or o.

3 Specifying Domains in LAO

In this section we address how to formally specify the domain in which an agent or robot is expected to live, in the language of LAO. Here, ϕ—with or without subscripts—denotes some (pre)condition expressed as a static sentence.

Firstly, axioms are required for action outcomes that say when an action is executable and for observations that say when (in which worlds) an observation is perceivable. A fundamental assumption in the reasoning about actions and change (RAC) community is that there must be one *executability axiom* for each action type. Executability axioms are similar to the precondition axioms in Reiter's situation calculus [12]. In multi-modal logic, one writes $\langle\alpha\rangle\top \leftrightarrow \phi$ to mean that it is possible to perform α if and only if the precondition ϕ holds. For instance, $\langle grab\rangle\top \leftrightarrow \neg holding$, $\langle drink\rangle\top \leftrightarrow holding \wedge full$, $\langle weigh\rangle\top \leftrightarrow holding$ and $\langle replace\rangle\top \leftrightarrow holding$ define in what worlds it is possible to execute each of the four available actions. There must be an executability axiom for each action in $\mathfrak{A}$.

We follow a systematic approach to specifying domain axioms that is based on the approach of Demolombe, Herzig and Varzinczak [3], which is in turn related to Reiter's approach with functions to systematize the specification of successor-state axioms in the situation calculus [12]. Here, there is only enough space to show the result of the systematic approach.

Effect axioms are required to capture the effects of actions. For the robot scenario, there exists an effect axiom $\neg holding \rightarrow (\langle grab\rangle holding \wedge \langle grab\rangle\neg holding)$

for the *grab* action. A translation for this axiom is, 'There exists an observation such that, if I am not holding the oil-can, it is possible that either I will be holding it or will still not be holding it, after grabbing it.' *grab* is a nondeterministic action with respect to *holding*. And the systematic approach produces the following effect axiom for *drink*: $full \land holding \rightarrow [drink]\neg full$.

Frame axioms and *condition closure axioms* state when actions do not have effects—these can easily be expressed in LAO [13].

In the same vein as executability axioms, we need *perceivability axioms*. However, to explain their specification, we define *ontic* (physical) actions and *sensory* actions. Ontic actions have intentional ontic effects, that is, effects on the environment that were the main intention of the agent. Sensory actions result in perception, and might only have (unintentional) side-effects.

Ontic actions (α_{ont}) each have a perceivability axiom of the form

$$(\forall o)\langle o \mid \alpha_{ont}\rangle\top \leftrightarrow o = obsNil.$$

For ontic actions, the null observation is perceived if and only if the action is executed. *grab*, *drink* and *replace* are ontic actions:

$$(\forall o)\langle o \mid grab\rangle\top \leftrightarrow o = obsNil;$$
$$(\forall o)\langle o \mid drink\rangle\top \leftrightarrow o = obsNil;$$
$$(\forall o)\langle o \mid replace\rangle\top \leftrightarrow o = obsNil.$$

For any instantiation of an observation o' other than *obsNil*, according to the semantics, $[o' \mid \alpha_{ont}]\bot$ is a logical consequence of these axioms.

Sensory actions typically have multiple observations and associated conditions for perceiving them. Sensory actions (α_{sen}) thus each have a set of perceivability axioms of the form

$$\langle o_1 \mid \alpha_{sen}\rangle\phi_1, \langle o_2 \mid \alpha_{sen}\rangle\phi_2, \ldots, \langle o_n \mid \alpha_{sen}\rangle\phi_n,$$

for stating when the associated observations are *possible*, where $\{o_1, o_2, \ldots, o_n\} = Dom(Q_{\alpha_{sen}})$[1] and the ϕ_i are the conditions. The following axioms state when observations associated with α_{sen} are *impossible*:

$$\neg\langle o_1 \mid \alpha_{sen}\rangle\neg\phi_1, \neg\langle o_2 \mid \alpha_{sen}\rangle\neg\phi_2, \ldots, \neg\langle o_n \mid \alpha_{sen}\rangle\neg\phi_n.$$

Note that the ϕ_i conditions need not characterize pair-wise disjoint sets of worlds, because more than one observation is allowed in the same worlds, given some action (see, e.g., *obsLight* and *obsMedium* below).

Lastly, to state that the observations not associated with action α_{sen} are always impossible given α_{sen} was executed, we need an axiom of the form

$$(\forall o)(o \neq o_1 \land o \neq o_2 \land \ldots \land o \neq o_n) \rightarrow \neg\langle o \mid \alpha_{sen}\rangle\top,$$

for each action.

[1] $Dom(Q_\alpha)$ is the set of all first elements in the pairs that make up Q_α.

Our only sensory action is *weigh* and its behavior with respect to perception can be captured by the following sentences.

$$\langle obsLight \mid weigh\rangle(\neg full \vee drank) \wedge [obsLight \mid weigh](\neg full \vee drank);$$
$$\langle obsHeavy \mid weigh\rangle(full \vee \neg drank) \wedge [obsHeavy \mid weigh](full \vee \neg drank);$$
$$\langle obsMedium \mid weigh\rangle\top \wedge [obsMedium \mid weigh]\top;$$
$$(\forall o)(o \neq obsHeavy \wedge o \neq obsLight \wedge o \neq obsMedium) \rightarrow \neg\langle o \mid weigh\rangle\top.$$

For instance, *obsLight* is perceivable given *weigh* was performed, if and only if either the oil-can is empty or the oil has been drunk.

All this can be done for deterministic and nondeterministic effects of actions [13]. All the axioms discussed in this section concern the dynamics of an environment. They are collectively the action laws and are here represented by LAW. All axioms in LAW are *global*, that is, true in every possible world. The state of affairs that an agent is in initially, can be characterized by a static, non-global sentence KB (knowledge base). The main task in LAO is to determine whether an arbitrary sentence φ is implied by KB, given LAW, that is, whether $LAW \models_G KB \rightarrow \varphi$. The next section shows how this can be done.

4 Tableaux for LAO

The tableau calculus we propose is adapted from Castilho, Gasquet and Herzig [1]. It is based on labeled formulae. It is a procedure to determine whether $\mathcal{K} \models_G \Psi$, where $\mathcal{K}$ is any set of global axioms in $\mathcal{L}_{LAO}$ and Ψ is any sentence in $\mathcal{L}_{LAO}$.

The tableau calculus for LAO, with all its rules, remarks and observations, is referred to as $\mathcal{C}_{LAO}$. The set of formulae to be checked (i.e., the initial set of formulae to which $\mathcal{C}_{LAO}$ must be applied) is called the *trunk*. A *labeled formula* is a pair (n, φ), where φ is a formula and n is an integer from the set of whole numbers, called the *label* of φ. A *skeleton* Σ is a ternary relation $\Sigma \subseteq (\Omega \cup \mathbb{N}) \times \mathfrak{A} \times \mathbb{N}$. Elements $(\cdot, a, n')$ of the relation are denoted $\cdot \xrightarrow{a} n'$. A *tree* $\mathcal{T}^i$ is a pair $\langle \Gamma^i, \Sigma^i \rangle$, where Γ^i is a set of labeled formulae and Σ^i is a skeleton. The initial tree is $\mathcal{T}^0 = \langle \{(0, \neg\Psi)\}, \emptyset \rangle$. Each $\mathcal{T}^{i+1}$ may be obtained from $\mathcal{T}^i$ by applying certain tableau rules to $\mathcal{T}^i$. Other rules add elements to Γ^i or Σ^i, producing a new state of the tree, but not necessarily a new tree.

Let a particular state of a tree be called a *node*. The application of a rule to a node k results in a new node k'. k' may be a node of the same tree as k, or k' may be the first node of a new tree. A *tableau* for the trunk is a set of trees $\mathcal{T}^0, \ldots, \mathcal{T}^n$ and their states, resulting from the application of *tableau rules* to the trunk and subsequent nodes. The tableau rules for LAO are:

- rule $\bot$: If Γ contains (n, φ) and $(n, \neg\varphi)$, then add $(n, \bot)$ to it.
- rule $\neg$: If Γ contains $(n, \neg\neg\varphi)$, then add (n, φ) to it.
- rule $\wedge$: If Γ contains $(n, \varphi \wedge \varphi')$, then add (n, φ) and (n, φ') to it.
- rule $\vee$: If Γ contains $(n, \neg(\varphi \wedge \varphi'))$, then add $(n, \neg\varphi)$ to it, and create $\mathcal{T}^i = \langle \Gamma \cup \{(n, \neg\varphi')\}, \Sigma \rangle$, where i is a new integer.

- rule =: If Γ contains $(n, c = c')$ and in fact, constants c and c' do not refer to the same constant, then add $(n, \bot)$ to it.
- rule $\forall$: If Γ contains $(n, (\forall v)\varphi)$ then add $(n, \varphi|^v_c)$ to it only if the constant c (of the right sort) appears in a formula in Γ.
- rule $\exists$: If Γ contains $(n, \neg(\forall v)\varphi)$ then add $(n, \neg\varphi|^v_{c_1} \vee \ldots \vee \neg\varphi|^v_{c_j})$ to Γ, for each constant in $\{c_1, \ldots c_j\}$ (of the right sort) that appears in the vocabulary.
- rule $\langle\alpha\rangle$: If Γ contains $(n, \neg[\alpha]\varphi)$, then add $(n, \neg(\forall o)[o \mid \alpha]\varphi)$ and $(n', \neg\varphi)$ to it, and add $n \xrightarrow{\alpha} n'$ to Σ, where n' is a fresh integer. For each $\beta \in \mathcal{K}$, add (n', β) to Γ.
- rule $[\alpha]$: If Γ contains $(n, [\alpha]\varphi)$ and Σ contains $n \xrightarrow{\alpha} n'$, add (n', φ) to Γ.
- rule $\langle o \mid \alpha\rangle$: If Γ contains $(n, \neg[o \mid \alpha]\varphi)$, then add $(n', \neg\varphi)$ to it, and add $n \xrightarrow{\alpha} n'$ and $o \xrightarrow{\alpha} n'$ to Σ, where n' is a fresh integer. For each $\beta \in \mathcal{K}$, add (n', β) to Γ.
- rule $[o \mid \alpha]$: If Γ contains $(n, [o \mid \alpha]\varphi)$ and Σ contains $n \xrightarrow{\alpha} n'$ and $o \xrightarrow{\alpha} n'$, then add (n', φ) to Γ.

The addition of $(n, \neg(\forall o)[o \mid \alpha]\varphi)$ to Γ in rule $\langle\alpha\rangle$ is due to Proposition 1. To make explicit that the formulae in $\mathcal{K}$ are global, they are all added to each new world (fresh integer) introduced in rules $\langle\alpha\rangle$ and $\langle o \mid \alpha\rangle$.

A tree $\langle\Gamma, \Sigma\rangle$ is *closed* if $(i, \bot) \in \Gamma$ for some i. It is *open* if it is not closed. A tableau is closed if all of its trees $\mathcal{T}^0, \ldots, \mathcal{T}^n$ are closed, else it is open.

Definition 4. *If a tableau for $\neg\Psi$ is closed (under $\mathcal{K}$), we write $\mathcal{K} \vdash_{LAO} \Psi$. If there is a* saturated *open tableau for $\neg\Psi$, we write $\mathcal{K} \nvdash_{LAO} \Psi$. A tableau is saturated if any rule that can be applied has been applied to all open trees.*

Theorem 1. *$\mathcal{C}_{LAO}$ is sound (if $\mathcal{K} \vdash_{LAO} \Psi$ then $\mathcal{K} \models_G \Psi$), complete (if $\mathcal{K} \models_G \Psi$ then $\mathcal{K} \vdash_{LAO} \Psi$) and decidable ($\mathcal{C}_{LAO}$ always terminates). [13]*

Using $\mathcal{C}_{LAO}$, the following can be proven:

- $LAW \models_G (full \wedge \neg drank \wedge \neg holding) \rightarrow \langle grab\rangle \neg holding$;
- $LAW \models_G (full \wedge \neg drank \wedge holding) \rightarrow (\exists o)[o \mid drink]\neg full$; and
- $LAW \models_G (full \wedge \neg drank \wedge holding) \rightarrow (\forall o)\langle o \mid drink\rangle \neg full$.

5 Discussion and Related Work

We believe that calculi based on first-order logic, like the situation calculus [9] and the event calculus [7] are too rich for our needs. We thus sought a simpler logic with the potential of being decidable. $\mathcal{LAP}$, the Logic of Actions and Plans, was found to be a suitable basis for our work. $\mathcal{LAP}$ is a multi-modal logic, close to but simpler than Propositional Dynamic Logic [5]. Castilho, Gasquet and Herzig [1] claim that it is sufficient to express most of the problems investigated in the field, however, it does not deal with sensing. To say that LAO is an extension of $\mathcal{LAP}$ is too strong. For example, their definition of $[\alpha]$ is the standard one for multi-modal logic, whereas the definition of $[\alpha]$ in LAO is not standard, in that in LAO its definition involves the *perceivability* relation Q_α.

A series of articles exists concerning probabilistic dynamic epistemic logic (PDEL) [6,14,15], which add probabilistic notions to dynamic epistemic logic (DEL) [16]. The language of PDEL includes formulae of the form $[A, e]\varphi$, where A is a "probabilistic update model" and e is an "event" from the domain of A. The terms *event* and *observation* often have the same meaning in probability theory. Observations in probability theory do not describe a state, but capture information about natural *occurrences*. The authors [15] allude that their events are closer to observations than to logical propositions. Therefore, as far as observations go, PDEL's $[A, e]\varphi$ corresponds to LAO's $[o \mid \alpha]\varphi$, however, the semantics of PDEL's operator is much richer. Note though, that PDEL is an epistemic logic, not a logic about action.

For our work, we have also found some inspiration from the language $\mathcal{ES}$ of Lakemeyer and Levesque [8], especially because $\mathcal{ES}$ has been extended to $\mathcal{ESP}$ [4] to include notions of probability—in a current line of investigation, we also intend to extend LAO with notions of probability. Although it is a situation-based logic, $\mathcal{ES}$ does not include situation terms. It is a second-order modal dialect with *object* and *action* sorts, and with universal quantification and equality. It has fluent and rigid functions and predicates. Fluent predicates include the special predicate *Poss* for defining preconditions on action executability, and a special predicate symbol for defining whether a sensing action was successful. The formula $[\alpha]\varphi$ in $\mathcal{ES}$ is defined to mean 'after α, φ is true'. The meaning is similar to that of the standard modal logic operator, although, in $\mathcal{ES}$, actions are deterministic. The diamond operator is not defined, but with *Poss* available, it needs not be defined. $\mathcal{ES}$'s $[\alpha]$ is thus also different to LAO's $[\alpha]$.

6 Concluding Remarks

Modal logic based RAC formalisms lack a straight-forward way to deal with sensing. In an attempt to solve the problem, we presented a multi-modal logic which includes reasoning about 'reified' observations on a par with actions. It was shown how to specify an agent domain in the language. We provided a tableau calculus ($\mathcal{C}_{LAO}$) for determining the validity of sentences of the logic, and it was stated that the calculus is sound, complete and decidable.

By adding observations to a simple dynamic logic explicitly, the resulting logic may be slightly more complex, while perhaps simplifying, for the domain expert, dealing with explicit observations (cf. Remark 1). Computational complexity of $\mathcal{C}_{LAO}$, and the influence of 'reifying' observations, must still be established. Since LAO is at least as expressive as multi-modal logic K and entailment here is global, we know that LAO is at least EXPTIME.

One of the main problems in systems for RAC is the *frame problem*. We have formulated a solution for LAO, which involves universal quantification over actions. For the interested reader, our *frame solution* appears in the accompanying technical report [13]. Alternatively, because LAO is essentially based on $\mathcal{LAP}$ one could in the future, adapt the frame solution from Castilho, Herzig and Varzinczak [2] to LAO.

Our next aim is to extend LAO to allow one to express uncertainty in action and perception by providing the machinery to specify *probabilistic* models (descriptions) for action outcomes and for perceivability of observations.

Acknowledgments

Thanks to Arina Britz, Gerhard Lakemeyer and the referees for their suggestions.

References

1. Castilho, M., Gasquet, O., Herzig, A.: Formalizing action and change in modal logic I: The frame problem. Journal of Logic and Computation 9(5), 701–735 (1999)
2. Castilho, M., Herzig, A., Varzinczak, I.: It depends on the context! A decidable logic of actions and plans based on a ternary dependence relation. In: 9th International Workshop on Nonmonotonic Reasoning, NMR 2002 (2002)
3. Demolombe, R., Herzig, A., Varzinczak, I.: Regression in modal logic. Journal of Applied Non-Classical Logics 13(2), 165–185 (2003)
4. Gabaldon, A., Lakemeyer, G.: $\mathcal{ESP}$: A logic of only-knowing, noisy sensing and acting. In: Proc. of 22nd Natl. Conf. on Artificial Intelligence (AAAI 2007), pp. 974–979. AAAI Press, Menlo Park (2007)
5. Harel, D., Kozen, D., Tiuryn, J.: Dynamic Logic. MIT Press, Cambridge (2000)
6. Kooi, B.: Probabilistic dynamic epistemic logic. Journal of Logic, Language and Information 12(4), 381–408 (2003)
7. Kowalski, R., Sergot, M.: A logic-based calculus of events. New Generation Computing 4, 67–95 (1986)
8. Lakemeyer, G., Levesque, H.: A semantic characterization of a useful fragment of the situation calculus with knowledge. In: Special issue in honor of John McCarthy, Artificial Intelligence, Elsevier, Amsterdam (2010)
9. McCarthy, J., Hayes, P.: Some philosophical problems from the standpoint of artificial intelligence. Machine Intelligence 4, 463–502 (1969)
10. Monahan, G.E.: A survey of partially observable Markov decision processes: Theory, models, and algorithms. Management Science 28(1), 1–16 (1982)
11. Popkorn, S.: First Steps in Modal Logic. Cambridge University Press, Cambridge (1994)
12. Reiter, R.: Knowledge in action: logical foundations for specifying and implementing dynamical systems. MIT Press, Massachusetts (2001)
13. Rens, G., Varzinczak, I., Meyer, T.: Partially observable Markov decision logic, Part I: Introducing explicit observation constants, KRR-10-01. Tech. rep., Knowledge Representation and Reasoning, Meraka Institute, CSIR, Pretoria, South Africa (June 2010), `http://krr.meraka.org.za/publications/2010`
14. Sack, J.: Extending probabilistic dynamic epistemic logic. Synthese 169, 124–257 (2009)
15. Van Benthem, J., Gerbrandy, J., Kooi, B.: Dynamic update with probabilities. Studia Logica 93(1), 67–96 (2009)
16. Van Ditmarsch, H., Van der Hoek, W., Kooi, B.: Dynamic Epistemic Logic. Springer, Dordrecht (2007)