

A Semantical Account of Progression in the Presence of Uncertainty

Vaishak Belle* and Gerhard Lakemeyer

Dept. of Computer Science

RWTH Aachen University

52056 Aachen

Germany

{belle,gerhard}@cs.rwth-aachen.de

Abstract

Building on a general theory of action by Reiter and his colleagues, Bacchus *et al.* give an account for formalizing degrees of belief and noisy actions in the situation calculus. Unfortunately, there is no clear solution to the projection problem for the formalism. And, while the model has epistemic features, it is not obvious what the agent's knowledge base should look like. Also, reasoning about uncertainty essentially resorts to second-order logic. In recent work, Gabaldon and Lakemeyer remedy these shortcomings somewhat, but here too the utility seems to be restricted to queries (with action operators) about the initial theory. In this paper, we propose a fresh amalgamation of a modal fragment of the situation calculus and uncertainty, where the idea will be to update the initial knowledge base, containing both ordinary and (certain kinds of) probabilistic beliefs, when noisy actions are performed. We show that the new semantics has the right properties, and study a special case where updating probabilistic beliefs is computable. Our ideas are closely related to the Lin and Reiter notion of progression.

1 Introduction

It is generally argued that when agents interact with an incompletely known world, in addition to reasoning about the effects of their actions, they need to be able to quantify their uncertainty about faulty sensors and effectors. Building on a logical theory of action by Reiter and his colleagues, Bacchus *et al.* [BHL] (1995) give an account for formalizing degrees of belief and noisy actions in the situation calculus. Their model can be thought of as two important extensions to the epistemic situation calculus (Reiter 2001): a set of axioms to capture nondeterminism in actions which still fall back on Reiter's solution to the frame problem, and a companion fluent to the epistemic one to capture a subjective assessment of uncertainty. Of course, a model of belief has limited appeal if it is not regularly updated with what is observed. BHL demonstrate that this is indeed the case.

Somewhat surprisingly, despite having epistemic features, it is not obvious what the agent's knowledge base actually looks like. There are also other problems. Since the situa-

tion calculus is defined axiomatically, reasoning about uncertainty, essentially a summation over the weights on epistemically accessible situations, resorts to second-order logic. For another, there is no clear solution to the *projection* problem, where we are to determine if a formula holds after a number of named actions have occurred, which is arguably a fundamental application of action formalisms.¹

To remedy a few of these shortcomings, Gabaldon and Lakemeyer [GL] (2007) extend a modal fragment of the situation calculus, called \mathcal{ES} , to include similar notions, which we review below. \mathcal{ES} also has epistemic features including only-knowing, which refers to all that the agent knows in the sense of having a knowledge base, and fully captures Reiter-style basic action theories allowing one to formulate successor state axioms of the form:

$$\begin{aligned} \forall v, y. \Box[v]wall = y &\equiv (v = adv1 \wedge wall = y + 1) \vee \\ &(v = reverse \wedge wall = y - 1) \vee \\ &v \neq adv1 \wedge v \neq reverse \wedge wall = y, \end{aligned}$$

which in English says that after any sequence of actions (\Box), doing an *adv1* brings the agent closer to the wall by a meter, and doing a *reverse* moves the agent away by a meter. Noisy or stochastic actions, which are to correspond to a number of ordinary actions such as *adv1* and *reverse*, are captured by including as part of the background theory:

$$\begin{aligned} \forall x, y. choice(move, x) = y &\equiv (x = adv1 \vee x = adv0) \\ &\wedge y = 1 \vee (x \neq adv1 \vee x \neq adv0) \wedge y = 0, \end{aligned}$$

$$prob(move, adv1) = .2, prob(move, adv0) = .8,$$

which in English says that executing *move* may correspond to doing *adv1* or doing *adv0*, with probabilities of .2 and .8 respectively. Now, if all that the agent knows is that the wall is 5 meters away and the above theory, then after doing a *move* it believes that it is 4 meters away with a probability of .2, as expected. The main technical leverage to BHL is that beliefs are evaluated semantically, avoiding the use of second-order representations.

While the results seem intuitive, it is also, unfortunately, not without problems. First, after doing noisy actions, they can only reason about probability formulas, whose semantics is quite involved. Moreover, here too the utility seems to be restricted to queries (with action operators) about the

*Vaishak Belle is supported by the B-IT graduate school and a DFG scholarship from the graduate school GK 643. Copyright © 2011, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

¹By *solution* we mean either the *regression* or the *progression* methodology (Reiter 2001). The latter is to be introduced shortly.

initial theory. Reasoning about actions this way is clearly unmanageable for anything but a small number of actions.

In a seminal paper, Lin and Reiter [LR] (1997) propose a powerful method to the projection problem called *progression*, where the idea is to update the initial theory. The advantage is clear: one can process multiple queries about the resulting state without any extra overhead. In recent work, Lakemeyer and Levesque [LL] (2009) propose a new semantics of \mathcal{ES} that allows us to use the progression methodology in the presence of only-knowing. The question arises if such a solution can also be lifted for a model of uncertainty.

In this paper, we propose a fresh amalgamation of \mathcal{ES} and uncertainty, where we show, among other things, what the knowledge base should look like after noisy actions. Our knowledge bases contain both ordinary beliefs, by which we mean first-order sentences taken to be the facts known by the agent, and probabilistic ones. We will propose a way where the ordinary beliefs can be progressed in a standard fashion, and probabilistic beliefs are updated in computable manner. To obtain these results, we will need to restrict the kinds of probabilistic beliefs in the knowledge base. Nevertheless, we believe the case we make is of practical interest. To the best of our knowledge, this is the first time that LR progression is investigated in the presence of uncertainty. For space reasons, this version of the paper only formulates results for noisy actions, leaving sensors for an extended version.

Organization is as follows. We begin with an introduction to the new logic. We then discuss basic action theories and our main results. Finally, we review related work and conclude. For reasons of space, this paper contains no proofs.

2 The Logic \mathcal{ES}_μ

Syntax Symbols are taken from the following vocabulary: first-order variables, second-order function variables, fluent and rigid functions, distinguished fluent *Poss*, rigid second-order functions, connectives $\neg, \forall, \wedge, [v], \llbracket v \rrbracket, \square, \mathbf{K}, \mathbf{B}, \mathbf{O}$.²

We assume that functions and variables come in three sorts: *object*, (*ordinary*) *action* and *stochastic action* with the understanding that actions are used with $[\cdot]$ and stochastic actions are used with $\llbracket \cdot \rrbracket$. Further, we assume that every fluent is of the object sort, and that distinguished symbols *prob* and *choice* are the only rigids of the object sort. Let

- \mathcal{N}_o be a countably infinite set of *object* names, e.g., *obj5*. It also includes the set of rational numbers \mathbb{Q} closed under standard arithmetical operations: $+, -, \times, \div$. Let $\mathbb{Q}_{[0,1]}$ denote the subset between 0 and 1 inclusive.
- $\mathcal{N}_a = \{A(m_1, \dots, m_k) \mid m_j \in \mathcal{N}_o, A \text{ is a function of the action sort}\}$ be action names, e.g., *adv1*, *drop(obj5)*.
- $\mathcal{N}_{st} = \{A(m_1, \dots, m_k) \mid m_j \in \mathcal{N}_o, A \text{ is a function of the stochastic action sort}\}$ be noisy action names, e.g., *move*.

Let \mathcal{N} be the set of all names, i.e., $\mathcal{N}_o \cup \mathcal{N}_a \cup \mathcal{N}_{st}$.

1. *Terms*: Every first-order variable and name is a term. We use \mathbf{t} to denote a vector of terms, and t_j to denote a term in \mathbf{t} . If f is a function and R is a second-order variable, then $f(\mathbf{t})$ and $R(\mathbf{t})$ are also terms.

²Predicates are excluded for simplicity. Also, other symbols, such as \vee, \exists, \equiv and \supset are understood as usual.

By *primitive term* and *primitive second-order term*, we mean ones of the form $f(\mathbf{m})$ and $R(\mathbf{m})$, where $m_j \in \mathcal{N}$.

2. *Formulas*: Suppose t, t' are terms. If α, β are formulas, then so are: $t = t', \alpha \wedge \beta, \neg\alpha, \forall x\alpha, \forall R\alpha, [t]\alpha, \llbracket t' \rrbracket\alpha, \square\alpha, \mathbf{K}\alpha, \mathbf{O}\alpha$, and $\mathbf{B}\alpha \geq r$ where $r \in \mathbb{Q}_{[0,1]}$.³

A *primitive formula* is of the form $f(\mathbf{m}) = n$, where $m_j \in \mathcal{N}$ and $n \in \mathcal{N}_o$. A *fluent* formula is one that does not mention *Poss, choice, prob, [v], \llbracket v \rrbracket, \square, \mathbf{K}, \mathbf{B} and \mathbf{O} . We also syntactically restrict formulas appearing in the scope of \mathbf{B} to be fluent formulas. We refer to formulas of the form $\mathbf{B}\alpha \geq r$ as *probability* or *belief* atoms.*

We write α_X^f to denote that every occurrence of function f in α is replaced by second-order variable X . We write α_x^t to mean that all free occurrences of variable x are substituted by term t . We read $[t]\alpha$ as “ α holds after action t ”, $\llbracket s \rrbracket\alpha$ as “ α holds after noisy action s ”, $\mathbf{K}\alpha$ as “ α is known”, $\mathbf{B}\alpha \geq r$ as “ α is believed with a probability $\geq r$ ”, and $\mathbf{O}\alpha$ as “only α is known”.⁴

Semantics Let \mathcal{Z} denote all finite sequences of names from \mathcal{N}_a , including $\langle \rangle$ (the empty sequence). The semantical apparatus is a possible-worlds framework, where *worlds* are functions from (non-rigid) primitive terms and \mathcal{Z} to \mathcal{N}_o , and from primitive second-order terms to \mathcal{N}_o . An *epistemic state* $e \subseteq \mathcal{W}$ is any set of worlds.

Terms are interpreted as follows. Names act as rigid designators, and the co-referring name for an arbitrary term is obtained wrt a world w as: $|t|_w = t$ if t is a name; $|f(\mathbf{t})|_w = w[f(\mathbf{n}), \langle \rangle]$ where $n_j = |t_j|_w$ and f is a function of the object sort; $|A(\mathbf{t})|_w = A(\mathbf{n})$ where $n_j = |t_j|_w$ as before and A is of the action or stochastic action sort; and $|R(\mathbf{t})|_w = w[R(\mathbf{n})]$.⁵

To interpret second-order variables, we introduce the notation $w \sim_X w'$ to mean that w and w' agree on everything except assignments involving X .

Next, to reason about noisy actions, we introduce functions Π and PR as follows. For every $t \in \mathcal{N}_a$, $\Pi(t) = t$ and for every $s \in \mathcal{N}_{st}$, $\Pi(s)$ is a finite set of action names. Finally, $\text{PR}(v) : \Pi(v) \rightarrow \mathbb{Q}_{[0,1]} - \{0\}$ is a distribution, i.e. choices obtain strictly positive probabilities.

Now, to interpret belief atoms over \mathcal{W} , we introduce the notion of a probability space (Halmos 1950).

Definition 1: A probability space is a tuple $(\mathcal{D}, \mathcal{X}, \mu)$ where \mathcal{D} is a set called the *sample space*, \mathcal{X} is a σ -algebra of subsets of \mathcal{D} (i.e. a set of subsets containing \mathcal{D} and closed under complementation and countable union), and a measure $\mu : \mathcal{X} \rightarrow [0, 1]$.

Known properties include: $\mu(\emptyset) = 0, \mu(\mathcal{D}) = 1$ and $\mu(A \cup B) = \mu(A) + \mu(B)$, if A and B are disjoint elements of \mathcal{X} . We only consider probability spaces that satisfy: $\mathcal{X} = 2^{\mathcal{D}}$, $\mu : \mathcal{X} \rightarrow \mathbb{Q}_{[0,1]}$, and $\mu(b) > 0$ for all $b \in \mathcal{X}, b \neq \emptyset$.

³While the only inequality operator we consider is \geq , this is wlog since expressing that α has a probability of r is easily written as $\mathbf{B}\alpha \geq r \wedge \mathbf{B}\neg\alpha \geq 1 - r$.

⁴The term “knowledge” is used with \mathbf{K} and “belief” with \mathbf{B} for readability purposes only. We do not insist that knowledge is true.

⁵We only need to obtain the co-referring name initially because the semantics we propose progresses worlds iteratively.

The idea is, much like (Fagin and Halpern 1994), to associate each world $w \in \mathcal{W}$ with a probability space. Suppose \mathfrak{S} is a function that maps worlds $w \in \mathcal{W}$ to a probability space $(\mathcal{D}^w, \mathcal{X}^w, \mu^w)$, where $\mathcal{D}^w \subseteq \mathcal{W}$. Fagin and Halpern (1994) discuss the various properties that result from how precisely \mathcal{D}^w is defined. Perhaps, the most intuitive model, is to let \mathcal{D}^w , for every $w \in \mathcal{W}$, be the epistemic state e . Uncertainty is then interpreted relative to the agent's beliefs, and this is what we will need. We simplify this idea further and only consider a *single* probability space over e .⁶

However, owing to our language, e may be uncountable. Instead of working with an infinite sample space (Halpern 2003), we proceed by reducing e to a (finite) set of equivalence classes of worlds. So let \mathcal{F} be a finite set of functions, H be a finite set of names, and $\mathcal{G} = \{f(m) \mid f \in \mathcal{F}, m_j \in H\}$. We write $w \approx w'$ to mean that for all $p \in \mathcal{G}$, $w[p, \langle \rangle] = w'[p, \langle \rangle]$. Now, given an epistemic state e , define $\|w\| = \{w' \in e \mid w' \approx w\}$ as the set of epistemically possible worlds that agree on \mathcal{G} initially.⁷ Analogously, for any $e' \subseteq e$, let $\|e'\| = \{\|w\| \mid w \in e'\}$ which is always finite. Intuitively, \mathcal{F} includes the fluents over which we define the basic action theory, and this sublanguage represents every fluent and name that an agent encounters during its operation. Putting this together, we define \mathfrak{S} as $(\|e\|, 2^{\|e\|}, \mu)$.

By a *model* we mean a tuple $\langle e, w, \mu, \delta \rangle$, where δ denotes the pair $\langle \Pi, \text{PR} \rangle$. The main purpose of our semantics is to clarify how fluent functions and belief atoms are to be understood. The account given is closely related to earlier work by LL, who define a notion of *progressing* epistemic states, and BHL, who give an account of how probabilities should be assigned to successor states. We briefly review both ideas:

- Given w and $t \in \mathcal{N}_a$, the *progressed world* w_t is defined as a world s.t. $w_t[p, z'] = w[p, t \cdot z']$ for all primitive terms p and actions sequences z' . Then, let $e_t = \{w'_t \mid w' \in e\}$.
- In the latter work, probabilities on a state are transferred to successor states when ordinary actions are performed, and are weakened by a factor on doing noisy ones.

So in a sense, the progression of worlds in $\|w\|$ must obtain the same probability as $\|w\|$ when ordinary actions are performed (analogously for noisy ones). Unfortunately, this would make the definition ill-defined in our case, mainly because even if w and w' are two different worlds, w_t and w'_t may be identical. This is not a bug. With progression we are essentially forgetting the past, but in BHL, the initial theory, and hence the initial situation(s), is kept around. It turns out, the only technical device we need is a notion of *normal worlds*, which fits very well with the purpose of the paper as we shall shortly see.

⁶As argued in GL, this often leads to agents holding precise beliefs about every formula. But consider a basket of oranges and bananas, where their proportion is not known. Then, we may not be able to assign an exact probability to the event "a selected fruit is orange". One remedy is to allow a set of measures to capture the entire range of possibilities. We ignore such issues for simplicity.

⁷Note that equivalence classes are understood wrt a particular epistemic state. But since it will be clear from the context which epistemic state we mean, we avoid the notational clutter.

Definition 2: Let w and w' be any two worlds, and suppose $w \approx w'$. They are said to be *normal* if $w_z \approx w'_z$ for all $z \in \mathcal{Z}|_H$, which is the restriction of \mathcal{Z} to all sequences that only mention names (of the action sort) from H . A *normal epistemic state* e is any set of normal worlds.

Definition 3: Given a model $M = \langle e, w, \mu, \delta \rangle$, where e is normal, its *progression* wrt $t \in \mathcal{N}_a$ is $\langle e_t, w_t, \mu_t, \delta \rangle$:

- let w_t and e_t be as above;
- for $w^* \in e_t$, let $\|w^*\| = \{w' \in e_t \mid w' \approx w^*\}$;
- let $\mu_t(\|w^*\|) = \mu(\cup_{\{w' \in e \mid w'_t \approx w^*\}} \|w'\|)$.

Definition 4: Given M as above, its *progression* wrt $s \in \mathcal{N}_{st}$, where $\Pi(s) = \{t_1, \dots, t_k\}$, is $\langle e_s, w_{t_j}, \mu_s, \delta \rangle$:

- let w_{t_j} and e_{t_j} be as above, and let $e_s = \cup e_{t_j}$;
- for $w^* \in e_s$, let $\|w^*\| = \{w' \in e_s \mid w' \approx w^*\}$;
- let $\mu_s(\|w^*\|) = \sum_j \mu(\cup_{\{w' \in e \mid w'_{t_j} \approx w^*\}} \|w'\|) \times \text{PR}(t_j)$.⁸

Before moving on, let us review what we have. Definition 2 says that if two (normal) worlds belong in an equivalence class, the same holds for their progressed versions. Definition 3 considers the possibility that when we progress worlds in e , different classes may merge since they end up agreeing on \mathcal{G} ; in which case, a sum of the weights on the earlier classes must apply to the merged one (to maintain normalization). Definition 4 follows the same principles, but this time for noisy actions defined over a distribution. One desirable property that we get from this machinery is that both $\mu_t(\|e_t\|)$ and $\mu_s(\|e_s\|)$ are always 1. For any $z = t_1 \cdot \dots \cdot t_k \in \mathcal{Z}$, we define $\langle e_z, w_z, \mu_z, \delta \rangle$ as the result of progressing $\langle e, w, \mu, \delta \rangle$ wrt z , in an iterative manner.

Given $\langle e, w, \mu, \delta \rangle$, the complete semantic definition is:

1. $e, w, \mu, \delta \models (t_1 = t_2)$ iff n_1 and n_2 are the same, where $n_j = |t_j|_w$;
2. $e, w, \mu, \delta \models \neg\alpha$ iff $e, w, \mu, \delta \not\models \alpha$;
3. $e, w, \mu, \delta \models \alpha \wedge \beta$ iff $e, w, \mu, \delta \models \alpha$ and $e, w, \mu, \delta \models \beta$;
4. $e, w, \mu, \delta \models \forall x\alpha$ iff $e, w, \mu, \delta \models \alpha_x^n$ for all names of appropriate sort;
5. $e, w, \mu, \delta \models \forall X\alpha$ iff $e, w', \mu, \delta \models \alpha$ for every $w' \sim_X w$;
6. $e, w, \mu, \delta \models [\tau]\alpha$ iff $e_t, w_t, \mu_t, \delta \models \alpha$, where $|\tau|_w = t$;
7. $e, w, \mu, \delta \models \llbracket \sigma \rrbracket \alpha$ iff $e_s, w_{t_j}, \mu_s, \delta \models \alpha$ for all $t_j \in \Pi(s)$, where $|\sigma|_w = s$;
8. $e, w, \mu, \delta \models \text{choice}(\sigma, \tau) = 1$ iff $t \in \Pi(s)$, where $|\sigma|_w = s$ and $|\tau|_w = t$;
9. $e, w, \mu, \delta \models \text{prob}(\sigma, \tau) = r$ iff $t \in \Pi(s)$ and $\text{PR}(t) = r$, where $|\sigma|_w = s$ and $|\tau|_w = t$;
10. $e, w, \mu, \delta \models \Box\alpha$ iff $e_z, w_z, \mu_z, \delta \models \alpha$ for all $z \in \mathcal{Z}$;
11. $e, w, \mu, \delta \models K\alpha$ iff for all $w' \in e$, $e, w', \mu, \delta \models \alpha$;
12. $e, w, \mu, \delta \models O\alpha$ iff for all $w', w' \in e$ iff $e, w', \mu, \delta \models \alpha$;
13. $e, w, \mu, \delta \models B\alpha \geq r$ iff $\mu(\llbracket \alpha \rrbracket_e) \geq r$;

⁸Here we mean $\text{PR}(s)[t_j]$, i.e., the probability assigned to t_j by the distribution $\text{PR}(s)$. We abbreviate this as $\text{PR}(t_j)$ for readability.

where, for a fluent formula α , $[\alpha]_e = \{w \mid w \models \alpha, w \in e\}$.

We say a sentence α is *true* for $\langle e, w, \mu, \delta \rangle$ if $e, w, \mu, \delta \models \alpha$. Given a set of sentences Σ , we write $\Sigma \models \alpha$ if for every normal e, w, μ, δ s.t. $e, w, \mu, \delta \models \alpha'$ for every $\alpha' \in \Sigma$, then $e, w, \mu, \delta \models \alpha$. Finally, we write $\models \alpha$ to mean $\{\} \models \alpha$.

For space reasons, we do not go over all of the general properties of the logic, but it suffices to say that \mathbf{K} satisfies the usual introspective properties of weak S5 (Hughes and Cresswell 1972). Additionally, we can demonstrate:

$$\begin{aligned} &\models \square(\mathbf{O}\alpha \supset \mathbf{K}\alpha); \\ &\models \square(\mathbf{K}\alpha \supset \mathbf{B}\alpha \geq r) \text{ for every } 0 \leq r \leq 1; \\ &\models \square(\mathbf{B}\alpha \geq r \supset \neg \mathbf{K}\neg\alpha), \text{ for every } 0 < r \leq 1; \\ &\models \square(\mathbf{B}(\alpha \wedge \beta) \geq r_1 \wedge \mathbf{B}(\alpha \wedge \neg\beta) \geq r_2) \supset \mathbf{B}\alpha \geq r_1 + r_2; \end{aligned}$$

The first property is about only-knowing and knowing (LL), and the remaining are about knowledge and probability, as in (Fagin and Halpern 1994), but where the properties hold in the presence of actions.

3 The Semantics of Progression

We begin by considering the equivalent of basic action theories (BATs) of the situation calculus.

Definition 5: Given a set of fluents \mathcal{F} , a set $\Sigma \subseteq \mathcal{E}\mathcal{S}_\mu$ is called a *basic action theory* over \mathcal{F} if it is the union of:⁹

1. the initial theory Σ_0 is any set of fluent sentences;
2. Σ_{pre} is a sentence $\square \text{Poss}(v) = 1 \equiv \pi$, where π is a fluent formula and $\text{Poss}(v) = 1$ denotes that v is possible;
3. Σ_{post} are sentences $\square[v]f(\mathbf{x}) = y \equiv \gamma_f(\mathbf{x}, y, v) \vee f(\mathbf{x}) = y \wedge \neg \exists h. \gamma_f(\mathbf{x}, h, v)$ where γ_f is a fluent formula, one for each f . We refer to them as successor state axioms (SSA).
4. Σ_{II} and Σ_{PR} are sentences of the form *choice*(x, x') = $y \equiv \psi$ and *prob*(x, x') = $y \equiv \varphi$ resp., where ψ and φ are fluent formulas only mentioning variables or names.

The idea is that Σ_0 expresses what is true initially, and there are sentences for the preconditions and effects of actions, which are formulated so as to incorporate Reiter's solution to the frame problem (Reiter 2001).¹⁰ The two additional components model stochastic actions. We often denote Σ_0 as ϕ , and the rest as $\square\beta$. We assume that $\phi \wedge \square\beta$ is all that an agent knows.¹¹ Finally, the agent's background theory also includes a conjunction of belief atoms, which we refer to as the *belief set*, and is handled separately from the above. We follow this terminology: by a *knowledge base* \mathcal{S} we mean a sentence of the form $\mathbf{O}(\phi \wedge \square\beta) \wedge \mathbf{B}\alpha \geq r$.

Example 1: Some aspects of a BAT are already dealt with in Section 1, where we have definitions for Σ_{post} , Σ_{II} and Σ_{PR} . In addition, let $\Sigma_{\text{pre}} = \{\square \text{Poss}(v) = 1 \equiv \text{true}\}$ (for simplicity), and let $\Sigma_0 = \{\text{wall} = 4 \vee \text{wall} = 5\}$ which the agent quantifies as say $\mathbf{B}(\text{wall} = 4) \geq .3 \wedge \mathbf{B}(\text{wall} = 5) \geq .7$ (denote by θ). Putting all this together, the agent's knowledge base is characterized as: $\mathbf{O}(\Sigma_0 \wedge \square\beta) \wedge \theta$.

⁹Free variables are assumed to be universally quantified.

¹⁰We remark that the unique name assumption for actions (Reiter 2001) is built into the logic.

¹¹Given any set of primitive formulas, an action theory determines precisely which of these are true after actions. Note that worlds that satisfy a BAT are *normal*. In this sense, e is normal.

In the sequel, we are concerned with the progression of our knowledge bases (\mathcal{S}). The question we must now answer is this: *what is progression in the presence of belief atoms and noisy actions?* Consider *classical progression*. In their work, LR give a model-theoretic definition for progressing basic action theories and discuss several properties that the new theory should satisfy. The main message is that for all queries about the future, the new and the initial theory behave identically (in a particular formal sense). It turns out an account similar to the LR notion also works for us.

Henceforth, let t (and s) denote a primitive action i.e., name, from \mathcal{N}_a (and name from \mathcal{N}_s resp.).

Definition 6: We call \mathcal{S}' the *progression* of \mathcal{S} wrt t (or s) iff for every model M , M is a model of \mathcal{S}' iff there is a model M' of \mathcal{S} s.t. M is the progression of M' wrt t (or s).

To get an idea about how one formulates \mathcal{S}' , let us consider the base theorem. In recent work, LL show that all that an agent knows after an (ordinary) action corresponds to the LR progression of an initial theory. When \mathbf{B} does not appear anywhere, we are able to reprove this result:

Theorem 1: If $\mathcal{S} = \mathbf{O}(\phi \wedge \square\beta)$, then its progression wrt t is $\mathbf{O}(\mathcal{P}(\phi) \wedge \square\beta)$ where:

$$\mathcal{P}(\phi) = \exists \mathcal{R} [\phi_{\mathcal{R}}^{\mathcal{F}} \wedge \bigwedge \forall \mathbf{x}, y. f(\mathbf{x}) = y \equiv \gamma_{t\mathcal{R}}^v], \mathcal{R} \text{ are second-order variables s.t. } R_j \text{ and } f_j \text{ are of the same arity.}$$

What this states is that if all that the agent knows is a BAT, then after t the agent knows another BAT, but where ϕ is replaced by $\mathcal{P}(\phi)$. While the general definition is second-order, under certain restrictions progression is first-order definable and efficient (Liu and Lakemeyer 2009; Belle and Lakemeyer 2011).

We now turn to the case when belief atoms do appear in \mathcal{S} . Fortunately, here too the only-knowing story is clear: we obtain a result on the set of sentences that the agent comes to know after ordinary and noisy actions. But it is not always clear what the belief set of \mathcal{S}' should be given an arbitrary \mathcal{S} . So unfortunately, we do not have proof as of yet if \mathcal{S}' exists in general. However, we are able to show that if it does exist, then it has the right properties, i.e. \mathcal{S}' is fully compatible with \mathcal{S} on unrestricted queries about the future.

Theorem 2: Suppose \mathcal{S}' is the progression of \mathcal{S} wrt t . Then, given any formula α , $\mathcal{S} \models [t]\alpha$ iff $\mathcal{S}' \models \alpha$.

(Analogously formulated for s .) For the rest of the paper, we are interested in a practical case, where \mathcal{S}' does exist and one obtains a representation theorem that characterizes \mathcal{S}' from \mathcal{S} in a computable manner.

3.1 Progression after ordinary actions

For the kind of applications we have in mind, it often suffices to maintain beliefs about (positive and negative) primitive formulas, as in the case of Example 1 and others in BHL. In general, we are interested in reasoning problems of the type:

Suppose wall = 5 is believed with a .7 probability. Then after doing a reverse, wall = 6 is now believed with the same probability.

Given this use case, it would be nice if after doing an action, a simple computational step, such as updating the value of a primitive term, is the only step needed (at least wrt beliefs). Unfortunately, once SSAs mention fluents about which the agent does not have complete information, then such a step is not possible. To this end, the only assumption we make when considering ordinary actions is that for all fluents mentioned in the belief set, denoted as $\mathcal{F}_B \subseteq \mathcal{F}$, we have (some sort of) complete knowledge about their SSAs.

Interestingly, expecting complete information about fluents in a SSA, even if only for belief atoms, is often too strong an assumption. To see why, consider Example 1. The problem here is that the new value of the fluent *wall* depends on the previous one, about which the agent is uncertain. In order to capture such useful cases, we introduce *normal SSAs*, where γ_f is a disjunction of formulas of the form:

$$\exists \mathbf{u}, b. [v = A(\mathbf{z}) \wedge \lambda_f(\mathbf{z}) \wedge f(\mathbf{x}) = b \wedge y = \alpha(b, \mathbf{z})]$$

where \mathbf{u} are the remaining variables from \mathbf{z} not appearing in \mathbf{x} , λ_f is any fluent formula not mentioning actions, and α is any arithmetical expression involving the previous value of $f(\mathbf{x})$ and \mathbf{z} .¹² So let SSAs for every $f \in \mathcal{F}_B$ be normal. Here is the complete information assumption that we are after:

Definition 7: A fluent sentence ϕ is *complete* wrt a set of primitive formulas \mathcal{A} if for all $l \in \mathcal{A}$, either $\phi \models l$ or $\phi \models \neg l$. It is complete wrt a fluent $f \in \mathcal{F}$, if it is complete wrt all instances of f which are primitive formulas.

Definition 8: We say $\mathcal{S} = \mathcal{O}(\phi \wedge \square\beta) \wedge \bigwedge \mathbf{B}l \geq r$ is *complete (for its belief set)* wrt an action theory if ϕ is complete wrt all fluents appearing in λ_f , for every $f \in \mathcal{F}_B$.¹³

Given a primitive action $A(e)$, denoted t , normal SSAs can be simplified, i.e., $\gamma_{f_t}^v$ simplifies to disjunctions of the form:

$$\exists b. [\lambda_f(e) \wedge f(\mathbf{x}) = b \wedge y = \alpha(b, e)].$$

By assumption, $\lambda_f(e)$ can be evaluated wrt the initial theory. Suppose $\phi \models \lambda_f(e)$. Then, the intuition is that, given any primitive formula of f that holds initially, say $f(\mathbf{m}) = n$, the instantiated SSA determines a precise value for $f(\mathbf{m})$ after t , i.e. $\alpha(n, e)$, which returns a name since α is an arithmetical expression. Of course, if $\phi \models \neg\lambda_f(e)$, then the value for $f(\mathbf{m})$ stays the same after doing t . In either case, if l denotes what holds initially, say $f(\mathbf{m}) = n$, then denote the formula $f(\mathbf{m}) = n'$, where n' is the value after t , as $\rho(l)$. We now present some results on updating belief sets:

Lemma 1: Let \mathcal{S} be complete (for its belief set), and l, l_j be primitive formulas. Given any model $\langle e, w, \mu, \delta \rangle$ of \mathcal{S} :

1. If $e, w, \mu, \delta \models \mathbf{B}l \geq r$ then $e_t, w_t, \mu_t, \delta \models \mathbf{B}\rho(l) \geq r$.
2. If $e, w, \mu, \delta \models \mathbf{B}l \geq r \wedge \mathbf{B}l \geq r'$ then $e_t, w_t, \mu_t, \delta \models \mathbf{B}\rho(l) \geq \max(r, r')$.
3. If $e, w, \mu, \delta \models \bigwedge \mathbf{B}l_j \geq r_j$, where l_j are different, then $e_t, w_t, \mu_t, \delta \models \bigwedge \mathbf{B}\rho(l_j) \geq \sum_{\{i|\rho(l_i) \text{ is the same as } \rho(l_j)\}} r_i$.

¹²Normal SSAs are strictly more general than *local-effects* (Liu and Lakemeyer 2009), where one constrains \mathbf{x} to be included in \mathbf{z} . Note: to express a SSA that does *not* depend on a previous value, let α simply not mention b .

¹³Complete knowledge about a SSA still allows for incomplete information. Here is an example: $\square[v]f(\mathbf{x}) = 1 \equiv v = \text{act}(\mathbf{x}) \wedge g = 1$, where (say) we know $g = 1$, but know nothing about f .

In other words, the new belief set is definable via simple steps. For illustration, showing (3), suppose $f(\mathbf{m}) = n$ and $f(\mathbf{m}) = n'$ are believed initially with weights r and r' . If t is s.t. it sets them both to $f(\mathbf{m}) = n^*$, then it is now believed with probability $r + r'$. For readability, we use $\mathbf{B}l \geq r$ in the following where we mean $\bigwedge \mathbf{B}l \geq r$, referring to Lemma 1 for the general case. Our representation theorem is this:

Theorem 3: If $\mathcal{S} = \mathcal{O}(\phi \wedge \square\beta) \wedge \mathbf{B}l \geq r$ is complete, then its progression wrt t is $\mathcal{O}(\mathcal{P}(\phi) \wedge \square\beta) \wedge \mathbf{B}\rho(l) \geq r$.

Example 1 continued. Suppose *reverse* is executed. The progression of Σ_0 is $\exists R. [(R = 4 \vee R = 5) \wedge \forall y. \text{wall} = y \equiv \text{reverse} = \text{reverse} \wedge R = y - 1]$, i.e. $\text{wall} = 5 \vee \text{wall} = 6$. Beliefs are updated similarly. Then, we get: $\mathcal{O}(\text{wall} = 5 \vee \text{wall} = 6 \wedge \square\beta) \wedge \mathbf{B}(\text{wall} = 5) \geq .3 \wedge \mathbf{B}(\text{wall} = 6) \geq .7$.

It follows that if t does not affect fluents in \mathcal{F}_B , then we can progress the initial theory while carrying over the belief set.

Theorem 4: Let \mathcal{S} be as above and t is s.t. for every $f \in \mathcal{F}_B$, $\Sigma_{\text{post}} \models [t]f(\mathbf{x}) = y \equiv f(\mathbf{x}) = y$. Then the progression of \mathcal{S} wrt t is $\mathcal{O}(\mathcal{P}(\phi) \wedge \square\beta) \wedge \mathbf{B}l \geq r$.

3.2 Progression after noisy actions

With noisy actions, a number of additional complexities arise. The subtlety here is that even where we have complete knowledge initially, doing a noisy action may result in non-trivial probabilistic beliefs in the progressed theory (see Section 1 for an example). To this end, we make a *local-effect* restriction (Liu and Lakemeyer 2009) with noisy actions.

We begin by grouping fluents that are affected when a noisy action is performed. Suppose the choices of $s \in \mathcal{N}_{st}$ are $A_1(e_1), \dots, A_k(e_k)$. Let \mathcal{F}_{A_j} denote the fluents f that A_j affects, i.e., A_j appears in the SSA of f . Let $\mathcal{F}_s = \bigcup \mathcal{F}_{A_j}$. Let H_s be the set of names in $\bigcup e_j$.

We now suppose that all fluents in \mathcal{F}_s have normal SSAs, where (additionally) the variables of an action (\mathbf{z}) mention all the variables appearing in the argument of the fluent (\mathbf{x}). The intuition is that every $A_j(e_j) \in \Pi(s)$ affects only a finite number of primitive terms $f(\mathbf{m}_1), \dots, f(\mathbf{m}_l)$ where $f \in \mathcal{F}_s$ and \mathbf{m}_i are names mentioned in e_j .

However, as hinted above, noisy actions generate probabilistic beliefs about fluents in \mathcal{F}_s . In this regard, we make the following reasonable assumption. Let a clause of the form $\bigvee f(\mathbf{m}) = n_j$ be called a *possible value* clause wrt $f(\mathbf{m})$. For efficiency reasons, we suppose that such clauses appear in the initial theory wrt all of $\{f(\mathbf{m}) \mid f \in \mathcal{F}_s, \mathbf{m}_j \in H_s\}$. We then assume that beliefs are maintained for each disjunct i.e. $f(\mathbf{m}) = n_j$. To see how this works, consider Example 1. Since the only fluent affected by *move* is *wall* and $\text{wall} = 4 \vee \text{wall} = 5$ is in the initial theory, the assumption essentially amounts to having beliefs about $\text{wall} = 4$ and $\text{wall} = 5$. In general, we do not believe the assumption is so serious because in most realistic domains, possible values range over a small number of names. And as is standard in probability literature, if the agent cannot specify a belief then all remaining possibilities are taken to be equally likely.

The above conditions have a clear reading when we think of fluents keeping information about the position of the robot

or a key combination (see BHL), where usually the fluents have no arguments, as is the case also with Example 1.

We now turn to our theorem. In what follows, we need to distinguish between the effects of the individual choices of s . To this end, let us denote by $\mathcal{P}(\phi, t_j)$ the progression of ϕ wrt ordinary action name $t_j \in \Pi(s)$, i.e. $\exists \mathcal{R} [\phi_{\mathcal{R}}^{\mathcal{F}} \wedge \wedge \forall x, y. f(x) = y \equiv \gamma_{t_j \mathcal{R}}^v]$. Analogously, let $\rho(l, t_j)$ denote the updated value of l when t_j is executed.

Theorem 5: *If $S = \mathbf{O}(\phi \wedge \square\beta) \wedge \mathbf{B}l \geq r$ is complete, then its progression wrt s is $\mathbf{O}(\bigvee \mathcal{P}(\phi, t_j) \wedge \square\beta) \wedge \theta$, where θ is $\wedge \mathbf{B}\rho(l, t_j) \geq r \times \sum_{\{i|\rho(l, t_i) \text{ is the same as } \rho(l, t_j)\}} \text{PR}(t_i)$.*

What this says is that since the agent does not know which choice was considered, the progression of Σ_0 wrt all choices are believed to be possible. In a sense, the agent ends up knowing less. This is what we would expect, and is also in the spirit of BHL. The belief section says that if one begins with a single atom, on doing a noisy action we obtain a conjunction of k atoms, where each atom is the result of updating wrt each of $\Pi(s)$. Now, if any of the resulting atoms are the same, then their probabilities will be added up. One sees that this is along the lines of what we have for ordinary actions, but where we now update wrt a set of actions.

Example 1 continued. On doing *move*, the progression of $wall = 4 \vee wall = 5$ is shown to be $wall = 3 \vee wall = 4 \vee wall = 5$. The new belief set is $\mathbf{B}(wall = 3) \geq .06 \wedge \mathbf{B}(wall = 4) \geq .38 \wedge \mathbf{B}(wall = 5) \geq .56$. So, owing to faulty execution, the belief in $wall = 5$ now is $.7 \times .8$ (*adv0* is executed) = $.56$. Quite analogously, the belief in $wall = 4$ is $.3 \times .8$ (in case *adv0* is executed given 4 meters) plus $.7 \times .2$ (in case *adv1* is executed given 5 meters) = $.38$. Belief in 3 meters is a result of *adv1* given 4 meters: $.3 \times .2$.

4 Related Work

Besides the situation calculus, a number of alternate formalisms for reasoning about action exist. Closely related include those based on the framework of dynamic logic (Harel, Kozen, and Tiuryn 2000), such as (Demolombe 2003), and the fluent calculus (Thielscher 2001). For a review on technical differences, see LL. Noisy actions are treated in BHL, in GL and by Thielscher (2001), but where LR progression is not considered. Progression, in fact, is not a new concept and is at the heart of most planning systems, starting with STRIPS. The most general account so far is by LR, where they also demonstrate how a number of special cases are compatible with their version.

The idea of assigning probabilities to possible worlds is inspired by earlier approaches (Fagin and Halpern 1994; Bacchus, Halpern, and Levesque 1995; Halpern 2003). Fagin and Halpern (1994) even consider the many agent case, but for a propositional language. They also do not explicitly address reasoning and actions. Halpern (2003) considers some first-order treatments of uncertainty, but he too does not address actions in that framework.¹⁴ Only-knowing is

¹⁴Some recent proposals deal with actions, knowledge and probability e.g., (Bentham, Gerbrandy, and Kooi 2009), but only for a propositional language.

also not considered in these proposals. A different treatment of only-knowing in the presence of uncertainty appears in GL, but where progression is not considered. Besides the area of KR, reasoning about probabilities is also of concern in game theory and in program verification; see (Fagin and Halpern 1994) for discussions.

5 Conclusions

The paper proposes a new model for reasoning about uncertainty and action, within a modal fragment of the situation calculus. Among its main features, is a semantics that clarifies LR progression in the presence of probabilistic beliefs and noisy actions. While we do not yet have a general result about the existence of progression in this setting, we did show that, when it does exist, it has the right semantic properties. We then presented a practical case, where we obtain representation theorems for characterizing the new knowledge base from the previous one. As an extension to this work, we also have results on noisy sensors, the details of which we leave to a longer version of the paper. We believe two future directions present themselves. It would be interesting to generalize the representation theorems for a larger class of action theories and for cases where we may express beliefs about quantified formulas. But perhaps the more pressing issue is to investigate if progression always exists, and if it does, whether it has a finite representation.

References

- Bacchus, F.; Halpern, J.; and Levesque, H. 1995. Reasoning about noisy sensors in the situation calculus. In *Proc. IJCAI*, 1933–1940.
- Belle, V., and Lakemeyer, G. 2011. On Progression and Query Evaluation in First-Order Knowledge Bases with Function Symbols. In *Proc. IJCAI*, to appear.
- Demolombe, R. 2003. Belief change: from situation calculus to modal logic. *Jour. of Appl. Non-Classical Logics* 13(2):187–198.
- Fagin, R., and Halpern, J. 1994. Reasoning about knowledge and probability. *J. ACM* 41(2):340–367.
- Gabalton, A., and Lakemeyer, G. 2007. ESP: A logic of only-knowing, noisy sensing and acting. In *AAAI*, 974–979.
- Halmos, P. 1950. Measure theory. *Van Nostrand Reinhold Company*.
- Halpern, J. 2003. *Reasoning about Uncertainty*. The MIT Press.
- Harel, D.; Kozen, D.; and Tiuryn, J. 2000. *Dynamic logic*. The MIT Press.
- Hughes, G. E., and Cresswell, M. J. 1972. *An introduction to modal logic*. Methuen London.
- Lakemeyer, G., and Levesque, H. 2009. A semantical account of progression in the presence of defaults. In *Proc. IJCAI*.
- Lin, F., and Reiter, R. 1997. How to progress a database. *Artificial Intelligence* 92(1-2):131–167.
- Liu, Y., and Lakemeyer, G. 2009. On first-order definability and computability of progression for local-effect actions and beyond. In *Proc. IJCAI*, 860–866.
- Reiter, R. 2001. *Knowledge in action: logical foundations for specifying and implementing dynamical systems*. The MIT Press.
- Thielscher, M. 2001. Planning with noisy actions (preliminary report). In *Proc. Australian Joint Conference on AI*.
- Van Benthem, J.; Gerbrandy, J.; and Kooi, B. 2009. Dynamic update with probabilities. *Studia Logica*, 93:1, 67–96.