Reasoning about Noisy Sensors in the Situation Calculus*

Fahiem Bacchus
Dept. Computer Science
University of Toronto
Toronto, Ontario
Canada, M5S 1A4
fbacchus@logos.uwaterloo.ca
416-978-7390, 416-978-1455 (FAX)

Joseph Y. Halpern
IBM Almaden Research Center
650 Harry Road
San Jose, CA 95120–6099
halpern@almaden.ibm.com

Hector Levesque
Dept. Computer Science
University of Toronto
Toronto, Ontario
Canada, M5S 1A4
hector@cs.toronto.edu
416-978-3618, 416-978-1455 (FAX)

Abstract: Agents interacting with an incompletely known dynamic world need to be able to reason about the effects of their actions, and to gain further information about that world using sensors of some sort. Unfortunately, sensor information is inherently noisy, and in general serves only to increase the agent's degree of confidence in various propositions. Building on a general logical theory of action formalized in the situation calculus, developed by Reiter and others, we propose a simple axiomatization of the effect on an agent's state of belief of taking a reading from a noisy sensor. By exploiting Reiter's solution to the frame problem, we automatically obtain that these sensor actions leave the rest of the world unaffected, and further, that non-sensor actions change the state of belief of the agent in appropriate ways.

Keywords: situation calculus, theories of action, knowledge, degree of belief.

Declaration: This paper has not already been accepted by and is not currently under review for a journal or another conference. Nor will it be submitted for such during IJCAI's review period.

*The work of Fahiem Bacchus and Hector Levesque was supported in part by the Canadian government through their NSERC and IRIS programs. Fahiem Bacchus is currently on sabbatical leave from the University of Waterloo, Ontario, Canada. Hector Levesque is a Fellow of the Canadian Institute for Advanced Research.

1 Introduction

Folk wisdom in AI used to have it that when it came to modeling perception and action (for robotics applications, for example), there was a fundamental decision that needed to be made, which determined to a large extent the character of the work to follow: Do we treat the uncertainty that results from noisy sensors and effectors quantitatively, using, for example, a *probabilistic* formalism such as Bayesian nets [Pea88], or do we use a *logical* model, and end up with formalisms such as STRIPS, the situation calculus, or some sort of modal logic? More recently, however, it is becoming increasingly clear that we can combine both probabilistic and logical reasoning in a single framework [Bac90, Hal90].

In this paper, we show how a logical account of perception and action, in particular, an account based on the situation calculus [MH69], can be augmented to deal quantitatively with the uncertainty that arises from noisy sensors. The type of reasoning we hope to capture includes the increase in confidence that is obtained by taking multiple sensor readings, and the progression of beliefs when ordinary actions (such as moving or grasping) are performed. There are standard probabilistic models that deal with the fusion of multiple sensor readings, and similarly, once a method is found for dealing with the frame problem, the situation calculus can deal quite well with ordinary actions. What interests us here is the integration of these notions.

An enriched version of the situation calculus, augmented with a solution to the frame problem proposed by Reiter [Rei91], has proved to be a very convenient formalism for modeling actions, their prerequisites, and effects. Although Reiter's proposal is limited in a number of ways, it has been extended to handle aspects of the ramification problem [LR94], agent ability [LLLS95], and continuous time [Pin94]. Another extension of the theory to deal with *complex actions* (sequence, iterations, concurrency, non-determinism, etc.), briefly described in Section 2, has led to a novel logic programming language called GOLOG. GOLOG has proven to be useful for describing high-level robot and softbot control [LLR95]. An implementation of GOLOG exists at the University of Toronto, and a number of small sample programs (including an elevator controller and a mail delivery robot) currently run in simulation mode. Whether or not the situation calculus will continue to be useful in a non-simulated robotic context, as additional extensions become necessary, or to what extent other representational formalisms can be put to similar use, remains to be seen. Nevertheless, by casting our work within this framework we hope to take advantage of these parallel developments.

Independently of the situation calculus, however, our formalism demonstrates an interesting interaction between ordinary actions and noisy perceptual actions. And as should be clear from our presentation, much of what we do here could be carried out in other logical frameworks.

The format of the rest of the paper is as follows. In the next section, we briefly review the theory of action in terms of which our account is formulated: the situation calculus, the solution to the frame problem proposed by Reiter [Rei91], and the extension, proposed by Scherl and Levesque [SL93], for dealing with knowledge. In Section 3, we consider how knowledge is affected by readings from noisy sensors. In Section 4, we augment the framework with probabilities, and present a simple formalization within the situation calculus of the degree of belief an agent has in propositions expressed as logical formulas. This allows us to formalize in more quantitative terms the changes in belief that arise from readings of noisy sensors. Examples of the formalism at work are presented in Section 5, and some conclusions are drawn in Section 6.

2 A Theory of Action

Our account of sensors is formulated as a logical theory T in an extended version of the situation calculus [MH69]. The situation calculus is a many-sorted dialect of the predicate calculus, containing sorts for (among other things) *situations*, which are like the possible worlds of modal logic; for primitive deterministic *actions*; and, since we will be dealing with probabilities, for *real numbers*. We assume the reader is familiar with the basic intuitions underlying the situation calculus; we briefly review the main ideas here.

2.1 The situation calculus and the frame problem

In this formalism, the world is taken to be in a certain state (or situation). Changes to the world arise only as the result of actions. This is modeled by having actions map situations to new situations using a special binary function symbol *do*. This function maps action-situation pairs to new situations, i.e., $s' = do(a, s)$ means that s' is the new situation that is the outcome of performing a in situation s. Predicates and functions whose values vary from situation to situation are called *fluents* and, by convention, take a situation as their last argument. We read, e.g., $P(\vec{x}, s)$ as "$\vec{x}$ has property P in situation s".[1]

The background theory T will contain axioms for the usual arithmetic operations on the real numbers, unique name axioms for actions, and various other foundational axioms for the situation calculus that need not concern us here. The domain dependent part of T consists of axioms characterizing the initial state of the world S_0, and the following: for every action type α, a *precondition axiom* of the form

$$\textsc{Poss}(\alpha, s) \equiv \Pi_\alpha(s),$$

where s is the only situation term mentioned in the formula $\Pi_\alpha(s)$; for every fluent F, a *successor state* axiom of the form

$$\textsc{Poss}(a, s) \supset [F(\vec{x}, do(a, s)) \equiv \Phi_F(\vec{x}, a, s)],$$

[1] Of course, in a modal logic, the possible worlds, s, are not part of the syntax, and we would write $s \models P(\vec{x})$ instead of $P(\vec{x}, s)$.

where s is the only situation term mentioned in the formula $\Phi_F(\vec{x}, a, s)$.[2]

For example, the precondition axiom for the drop action might assert that it is possible for the agent to drop an object x in situation s iff the agent is holding x in s: $\text{POSS}(\mathsf{drop}(x), s) \equiv \textit{Holding}(x, s)$. For the fluent *Broken*, a successor state axiom might assert that x is broken after the action a iff x was fragile and the agent dropped it, or x was broken and the agent did not repair it: $\text{POSS}(a, s) \supset \textit{Broken}(x, \textit{do}(a, s)) \equiv (a = \mathsf{drop}(x) \wedge \textit{Fragile}(x, s)) \vee (\textit{Broken}(x, s) \wedge a \neq \mathsf{repair}(x))$.

These axioms incorporate a treatment of the classic frame problem [MH69] proposed by Reiter [Rei91], extending previous proposals by Pednault [Ped89], Schubert [Sch90] and Haas [Haa87]. In particular, Reiter shows how the successor state axioms above can be automatically generated from a collection of simple effect axioms describing only the *changes* that result from performing an action. Frame axioms need not be enumerated since they are entailments of the successor state axioms.[3]

Reiter's solution to the frame problem applies only to primitive deterministic actions. However, Levesque et al. [LLR95] show how, as in dynamic logic [Har79], primitive actions can be composed in various ways to generate an expressive class of complex actions. Specifically, they show that there is a situation calculus formula, which we abbreviate by $\textit{Do}(a, s, s')$, that expresses the proposition that s' is one of the possible outcomes of doing complex action a starting in situation s. Here we only need one type of complex action: the nondeterministic choice of an action from a parameterized family of actions. Let $a(x)$ be a family of primitive actions parameterized by x. For example, a might be the action "approach the wall" and x might be a numeric parameter specifying the distance to be moved. The complex action $(\pi x).a$, can be read as "perform primitive action $a(x)$ for some nondeterministically selected value of x". In this case, $\textit{Do}((\pi x).a, s, s')$ stands for $\exists x.\text{POSS}(a(x), s) \wedge s' = \textit{do}(a(x), s)$. Note that since complex actions ultimately reduce to primitive ones, their preconditions, effects and non-effects are automatically entailed.

2.2 Knowledge and action

Scherl and Levesque [SL93] provide another extension to Reiter's basic approach by incorporating an epistemic state for the agent. To characterize this epistemic state in the language of the situation calculus, they follow Moore [Moo85] and introduce a new binary fluent K. The K fluent acts as a binary relation on situations, in an analogue of the accessibility relation among possible worlds in modal logics. Intuitively, $K(s', s)$ holds if in situation s, the agent considers the situation s' to be possible.[4] As in modal logic, knowledge is defined as truth in all accessible situations. Here, $\text{KNOW}(\phi, s)$ is an abbreviation for the formula $\forall s'.K(s', s) \supset \phi[s']$, where we assume that the situation argument has been removed from the fluents in

[2]This is the axiom for predicate fluents; the axiom for functional fluents would be analogous.

[3]Reiter's solution ignores the ramification problem; a treatment compatible with the approach has been proposed by Lin and Reiter [LR94].

[4]For compatibility with other fluents, the order of arguments is the opposite of the standard reading in modal logic.

ϕ and $\phi[s']$ is the result of introducing s' as a new situation argument. Thus, for example, $\text{KNOW}(\neg \mathit{Broken}(x), s)$ is an abbreviation for $\forall s'. K(s', s) \supset \neg \mathit{Broken}(x, s')$. For simplicity, we take K to be transitive and Euclidean, which ensures that the agent always knows whether or not it knows something (i.e., the agent has the power of positive and negative introspection).

To characterize how knowledge is affected by actions while preserving Reiter's solution to the frame problem, Scherl and Levesque first assume that all primitive actions can be divided into ordinary actions, such as drop and repair, and special knowledge-producing actions that affect K alone. For example, the agent might have available an action exactRead, whose effect is to change the agent's knowledge state so that it comes to know the exact distance to the wall in front of it. Scherl and Levesque assume that after performing an action, of either type, the agent will know that he performed it, and thus that the action was possible. For knowledge-producing actions like exactRead, the agent will also come to know the value of some fluent associated with the action, like *wallDist*, the distance to the wall. If exactRead is the only knowledge-producing action, we end up with the following successor state axiom for K:[5]

$$\begin{aligned} &\text{POSS}(a, s) \supset K(s', \mathit{do}(a, s)) \equiv \\ &\quad \exists s''. s' = \mathit{do}(a, s'') \wedge K(s'', s) \wedge \text{POSS}(a, s'') \wedge \\ &\qquad a = \mathsf{exactRead} \supset \mathit{wallDist}(s'') = \mathit{wallDist}(s) \end{aligned} \tag{1}$$

This entails $\text{POSS}(\mathsf{exactRead}, s) \supset \exists d. \text{KNOW}(d = \mathit{wallDist}, \mathit{do}(\mathsf{exactRead}, s))$: after doing exactRead, the agent knows the distance to the wall. To see this consider the situations that are K-related to $\mathit{do}(\mathsf{exactRead}, s)$, the successor of s. All such situations s' have the property that they are the successor states of some other situation s'' in which the distance to the wall is the same as it is in s. Since, exactRead does not change the distance to the wall, the successor state axiom for *wallDist* ensures that $\mathit{wallDist}(s') = \mathit{wallDist}(s'')$. Hence, all of the situations K-related to $\mathit{do}(\mathsf{exactRead}, s)$ have the same value for this fluent, and our observation follows.

3 Sensors and Noise

One problem with the Scherl and Levesque account, is that it is unrealistic to assume that an agent has available an exactRead action that allows it to learn the *exact* distance to the wall. A more realistic assumption is that the agent is in possession of a number of *sensors*, that give it some information about, but not exact knowledge of, various fluents. We expect a sensor reading to be correlated with, but not a deterministic function of, the quantity being measured. For example, we might imagine that there is a sonar sensor that can be used to measure the distance to the nearest wall. There might also be a laser range finder used to measure the distance to the wall, but it might be correlated with the actual distance in a different way.

[5] With additional knowledge-producing actions additional clauses are introduced into K's successor state axiom, one clause for each knowledge-producing action. These clauses will be very similar to that given for $a = \mathsf{exactRead}$.

There are various ways of modeling this. We present one here, motivated by our desire to have the basic actions be deterministic (and thus preserve the simple solution to the frame problem). Assume we have an action of the form $\mathsf{observe}(x)$, that occurs whenever the agent observes reading x on the sonar. If we assume that the sonar reading is always within b units of the true distance to the wall (rather than being equal to the distance to the wall, as in the previous example), then we get the following precondition axiom:[6]

$$\text{POSS}(\mathsf{observe}(x), s) \equiv |\mathit{wallDist}(s) - x| \leq b.$$

If we now assume, as did Scherl and Levesque, that an agent learns that an action is possible by successfully performing it, it will follow that after an observe action, the agent will learn the distance to the wall to within b units. In other words, the Scherl and Levesque successor state axiom for K from the previous section entails $\text{POSS}(\mathsf{observe}(x), s) \supset \text{KNOW}(|\mathit{wallDist} - x| \leq b, \mathit{do}(\mathsf{observe}(x), s))$, by an argument analogous to the one for exactRead, but now using the POSS predicate. In this case, with a precondition axiom as above, it is not necessary to treat observe, or similar observation actions from other sensors, as a special knowledge-producing action; it is instead treated like ordinary actions such as drop and repair.

Of course it is somewhat odd to say that the *agent* performs an action such as $\mathsf{observe}(3.7)$, as if it had the choice of performing, say, $\mathsf{observe}(3.6)$ instead. What we would prefer to say is that the agent decides to read the sonar, and that what *happens*, is that 3.7 is observed.

This can be modeled by using a nondeterministic composition of the primitive $\mathsf{observe}(x)$ actions. We define a complex action read as follows:

$$\mathsf{read} \stackrel{\text{def}}{=} (\pi x).\mathsf{observe}(x).$$

Given the abbreviation *Do* defined above, this means that

$$\mathit{Do}(\mathsf{read}, s, s') \equiv \exists x. |\mathit{wallDist}(s) - x| \leq b \wedge s' = \mathit{do}(\mathsf{observe}(x), s).$$

Using the successor state axiom for K, we get the following:

$$\mathit{Do}(\mathsf{read}, s, s') \supset \exists x. \text{KNOW}(|\mathit{wallDist} - x| \leq b, s').$$

So reading the sonar in s entails getting to a state where the agent has observed a (non-deterministically selected) consistent sonar value x. Moreover, the agent knows in that state the appropriate bound on the true distance to the wall. It is easy to check that doing several consecutive sonar readings can increase the agent's knowledge about the true distance to the wall (i.e., tighten the interval that the agent knows contains the true distance to the wall) and never decrease it. Similar considerations apply to other sensors whose read actions would be defined analogously.

[6]This particular precondition axiom only mentions the error bound, but other conditions can be included here as well.

4 Probability

Suppose we have a sensor with an error bound of $b = 2$, and we make a number of readings of a particular fluent using the sensor, all of which are clustered around the value 3. For concreteness, suppose they are all between 2.8 and 3.1. As far as *knowledge* goes, all the agent will be able to conclude is that he knows the fluent to have a value in the range [1.1,4.8]. Getting numerous readings of 3 will not change this knowledge. Yet, even if the agent is using a cheap sensor, we might hope that getting such readings would increase the agent's *degree of belief* that the true value of the fluent is very close to 3.

To formalize these intuitions, we introduce a probability distribution over the agent's set of K-related states. In particular, we associate with each situation in this set a relative weight. Intuitively, the relative weight measures the degree to which the agent believes that situation is in fact the real situation. However, it is more convenient to avoid forcing this weight to be a probability; instead we only require that these weights be positive and that their sum over all of the K-related states be finite. To obtain a true probability, we will simply normalize these weights so that they do in fact sum to 1.

Syntactically, we introduce a new functional fluent $p(s', s)$ whose value is the weight the agent assigns to situation s' when it is in situation s. This weight is unnormalized, and we introduce an abbreviation $\text{BEL}(\phi, s)$ to refer to the agent's *probabilistic* degrees of belief. Specifically, $\text{BEL}(\phi, s)$ is a number from 0 to 1 that is intended to stand for the agent's degree of belief in the assertion expressed by ϕ, when it is in situation s. As with KNOW, the first argument to BEL will be a formula containing fluents that are missing a situation argument, and we use the notation $\phi[s']$ as before for the formula that results when s' is introduced as the new situation argument. Informally, $\text{BEL}(\phi, s)$ will be defined to be the sum of the p weights of the accessible situations where ϕ holds, divided by the sum of the p weights of all accessible situations:

$$\sum_{s':K(s',s)\,\wedge\,\phi[s']} p(s', s) \bigg/ \sum_{s':K(s',s)} p(s', s).$$

These summations can be formalized within the situation calculus. In conjunction with Equation 2, below, the logical consequence of this formalization is that BEL is a probability distribution. The details of this development will be provided in a latter full report on this work.[7] For now, what matters is that we have something of the form

$$\text{BEL}(\phi, s) = x \stackrel{\text{def}}{=} \ldots \quad \textit{formula involving } p \quad \ldots$$

and that $\text{BEL}(\bullet, s)$ is a probability distribution over the situations K-related to s.

To ensure that the normalized sums of p, i.e., BEL, are in fact probabilities requires a constraint on the values of p in the initial state S_0. The following constraint

[7]Technically, we are assuming that we have discrete probability functions. This means, for example, that only countably many possible distances to the wall have probability greater than 0.

must be added to the background theory T:

$$\forall s. K(s, S_0) \supset [p(s, S_0) \geq 0 \wedge \forall s'. p(s', s) \geq 0]. \tag{2}$$

Since *p* is a fluent, we need to say how it is affected by actions. As with our treatment of every other fluent, we want to develop a successor state axiom for *p*. Many actions will have only an indirect effect on the agent's beliefs; the agent will only come to know that the action was successfully performed and this will affect its beliefs about the fluents changed by the action. For such actions, we want

$$p(s', do(a, s)) = \textbf{ If } s' = do(a, s'') \textbf{ then } p(s'', s) \textbf{ else } 0.$$[8]

This simply projects the relative degree of belief in s'' to its successor s'.

Notice that in making the projection we are transferring the agent's beliefs to situations with different properties. (This is somewhat related to Lewis's notion of *imaging* [Lew76]). In these new situations, all of the changes due to action a have occurred. For example, say that $\mathsf{approachW}(x)$ is the action "move x units towards the wall". In this case, the above equation will imply

$$\text{BEL}(\mathit{wallDist} = z, \mathit{do}(\mathsf{approachW}(x), s)) = \text{BEL}(\mathit{wallDist} = z + x, s)$$

Thus, if the agent believed it highly likely that she was 10 units from the wall in situation s, then she would believe it just as likely that she was 9 units from the wall after moving towards the wall 1 unit.

Things are a little more complicated when we have to deal with primitive actions like $\mathsf{observe}(x)$. As we mentioned before, we do not really think of this as an action that the agent performs; the agent is actually performing the read action. Although we have modeled read as a nondeterministic choice among $\mathsf{observe}(x)$ actions, it is actually better thought of as a probabilistic choice. Moreover, the probability of getting x as the reading depends on the situation and the accuracy of the sensor. In the simplest case, we would expect that in situation s, the smaller $|\mathit{wallDist}(s) - x|$, the greater the probability of $\mathsf{observe}(x)$; the exact distribution, however, will depend on the sensor.

To make this precise, we propose that for every sensor i, there is a likelihood function ℓ_i, where $\ell_i(x, s)$ denotes the probability of obtaining a reading of x from sensor i in situation s. Different applications will want to characterize these likelihood functions values differently, dependent on how complicated a model of sensor error is desired; here we simply assume that for each sensor i, the background theory T contains a *sensor noise axiom* of the form

$$\ell_i(x, s) = \Gamma_i(x, s),$$

where $\Gamma_i(x, s)$ is a term whose value is always between 0 and 1, and is equal to 0 when x exceeds the error bounds of the sensor (if there are any error bounds).

[8] This, of course, is an abbreviation for the formula $p(s', do(a, s)) = y \equiv (\neg \exists s''. s' = do(a, s'') \wedge y = 0) \vee (\exists s''. s' = do(a, s'') \wedge y = p(s'', s))$. We continue to use such abbreviations below.

For example, we might want to say that the likelihood of getting a sonar reading of x depends only on the difference between x and the current *wallDist*, and that this difference, i.e., the sonar's noise, is normally distributed with mean 0 and standard deviation σ. In this case, we would have an axiom of the form

$$\ell_{sonar}(x, s) = \mathit{Normal}\left(\frac{\mathit{wallDist}(s) - x}{\sigma}\right),$$

where $\mathit{Normal}(z)$ is a (discrete version of) the normal density function with mean 0 and standard deviation 1. This function could be defined in T by a simple table of values.[9]

Given such a function, for a situation $s' = do(\mathsf{observe}(x), s'')$ accessible from $do(\mathsf{observe}(x), s)$, we want to weigh the degree of belief in s' by $\ell_{sonar}(x, s'')$. That is, we want

$$\begin{aligned} p(s', do(\mathsf{observe}(x), s)) = \quad & \textbf{If } s' = do(\mathsf{observe}(x), s'') \\ & \textbf{then } p(s'', s) \times \ell_{sonar}(x, s'') \\ & \textbf{else } 0. \end{aligned}$$

To get this property for $\mathsf{observe}$ actions, and the one above for ordinary actions, we use the following general successor state axiom for *p*, which we include as part of the background theory T:

$$\begin{aligned} p(s', do(a, s)) = \quad & \textbf{If } s' = do(a, s'') \\ & \textbf{then } p(s'', s) \times L(a, s'') \\ & \textbf{else } 0, \end{aligned} \tag{3}$$

where

$$\begin{aligned} L(a, s) \stackrel{\text{def}}{=} \quad & \textbf{If } a = \mathsf{observe}_1(x) \textbf{ then } \ell_1(x, s) \textbf{ else} \\ & \quad \vdots \\ & \textbf{If } a = \mathsf{observe}_k(x) \textbf{ then } \ell_k(x, s) \\ & \textbf{else } 1 \end{aligned}$$

This completes our formal characterization of adding probability to the situation calculus. So apart from the abbreviations noted above, we have exactly 3 situation calculus axioms: the Scherl and Levesque successor-state axiom for K, a constraint on *p* in the initial state, and a successor-state axiom for *p*.

5 Properties of the Formalization

Our formalization of noisy sensors in the situation calculus is extremely simple, yet it has some very interesting consequences. In this section we explore some of these consequences, mostly through example. Before turning to the examples, however,

[9] Here we are using the standard transformation: $\mathit{Normal}((z + m)/\sigma)$ is the density function of a normal distribution with mean m and standard deviation σ.

we discuss in more detail one aspect of our approach that might at first glance seem puzzling.

We have modeled the agent reading its sensors as the execution of read, a nondeterministic choice among observe(x) actions. It would seem to be more appropriate to model read as a probabilistic action. And, in particular, it is not immediately obvious that a nondeterministic action is sufficient for our purpose.

Probabilistic actions were used in an approach explored by Halpern and Tuttle [HT93]. There, a general model of knowledge and probability in systems is provided, based on the framework of [FHMV95]. Roughly speaking, a *system* in [HT93] is identified with a set of *runs* or *executions*, where a run is a function from the natural numbers to situations. In the context of the situation calculus, we can think of a run as a sequence $s_0, s_1, s_2, \ldots$ of situations, such that $s_{i+1} = do(a, s_i)$ for some primitive action a. In [HT93], the actions were allowed to be probabilistic, not just nondeterministic. Probabilistic actions induce probabilities over the runs. From the probabilities on runs, we get, at each situation s, a probability on other situations, by conditioning on what the agent knows at that situation.

Here we also have probabilities at situations, but instead of being generated directly by the probabilities of actions, our probabilities are generated by our treatment of likelihoods in the definition of the successor-state axiom for *p*. Nevertheless, we can show that, at each situation, the agent places the same probabilities on other situations in the model we provide here as in the corresponding model of [HT93]. In the full version of the paper, we will show that for all formulas in our language, we get the same answer as we would using the framework of [HT93]. This shows that in our context non-deterministic read actions are sufficient, and that they yield the same answers that one would get from probabilistic read actions.

Belief Update. One of the standard probabilistic models of sensors is to assume that we have two pieces of probabilistic information: a prior distribution $\Pr(t)$ on the value t being sensed, and a conditional distribution $\Pr(x|t)$ that gives us the probability of sensing x given that the true value is t. Furthermore, the standard model requires the assumption that the value read from the sensor is only dependent on the true value, and is thus independent of other factors given this value.

We can now apply Bayes Rule to obtain a posterior probability $\Pr(t|x)$ over the values t given that the sensor read the value x: $\Pr(t|x) = \Pr(x|t)\Pr(t)/\Pr(x)$. The denominator is the only expression we do not know, but this is just a normalizing factor equal to $\sum_t' \Pr(x|t')\Pr(t')$. The key factor is the numerator $\Pr(x|t)\Pr(t)$ that describes the relative probability of different values of t given the observation x.

If we make a similar set of assumptions in our framework we obtain exactly this probabilistic model of the effect of sensing on the agent's beliefs. Let I be a set of sensor noise axioms of the form $\ell_i(x, s) = Err_i(x, e_i(s))$, where e_i is the fluent that is sensed by sensor i, and $Err_i(x, e)$ is some expression with just two free variables (both numeric), x and e. A sensor noise axiom of this form captures, in the language of the situation calculus, the assumption that the probability of obtaining a reading of x from sensor i in situation s, i.e., $\ell_i(x, s)$, is dependent only on the true value of the fluent being sensed. That is, no other properties of s affect this probability.

Proposition 1 *Let T be the background theory that includes the axioms given in Eq. 1–3. Then*

$$T \cup I \models \qquad (4)$$
$$\text{BEL}(e_i = t, do(\mathsf{observe}_i(x), s)) = \frac{\text{BEL}(e_i = t, s) Err_i(x, e_i(s))}{\sum_{t'} \text{BEL}(e_i = t', s) Err_i(x, e_i(s))},$$

where the denominator is simply a normalizing factor.

If the sensor is informative, i.e., if it is more likely to read values closer to the true value than values further away, then this proposition ensures that the agent's beliefs about the fluent he is sensing will become sharpened about the sensed value.

Example 5.1: Say that the agent is sensing the distance to the wall, *wallDist*, using a read of its sonar sensor. Let $\ell_{sonar}(x, s) = Err(x - \textit{wallDist}(s))$. That is, not only do we assume that the sonar's error is dependent only on the true value of the fluent being sensed (assumption I above), but we also assume that this error is characterized by a simple additive noise model. The sonar reads the true value plus a noise component. Hence, the probability of obtaining the reading x given that the true value is *wallDist*(s) is a function of the difference between the two (i.e., a function of the noise). For definiteness, let $Err(0) = 0.5$, $Err(-1) = 0.25$, and $Err(1) = 0.25$. (The probability is zero that the sonar will read a value that is more than 1 unit away from the true value). Let the agent's beliefs about *wallDist* in S_0, for $y = 11$, 12, 8, and 9, be $\text{BEL}(\textit{wallDist} = y, S_0) = 1/8$, and $\text{BEL}(\textit{wallDist} = 10, S_0) = 1/2$. (The agent has zero belief in the distance being any other value). This distribution of beliefs for the various values of *wallDist* in S_0 are shown in Figure 1.

Suppose that the agent reads its sonar, and observes the value 11. In the new situation $S_1 = do(\mathsf{observe}(11), S_0)$, a simple calculation using the above formula shows how the agent's beliefs have been altered. The new distribution is shown in the figure. Since the sonar has probability zero in being more than 1 unit away from the true value, the agent now has zero degree of belief in the values 8 and 9.[10]

Note that the diagram shows that the agent still believes that *wallDist* $= 10$ is the most likely value, even though its sonar returned the value 11. This arises from the agent's high prior belief in *wallDist* $= 10$.

Sequences of sensor readings of the same fluent, including sequences of readings from different sensors, are also handled correctly in our framework. Such sequences

[10] If the agent *knows* these error bounds, i.e., if these bounds are part of the preconditions for the sonar, it will come to know that *wallDist* is in the range 10–11. On the other hand, if the agent only has zero degree of belief in these outcomes, it will come to believe with degree 1 that *wallDist* is in that range. That is, our framework can distinguish between full belief and knowledge.

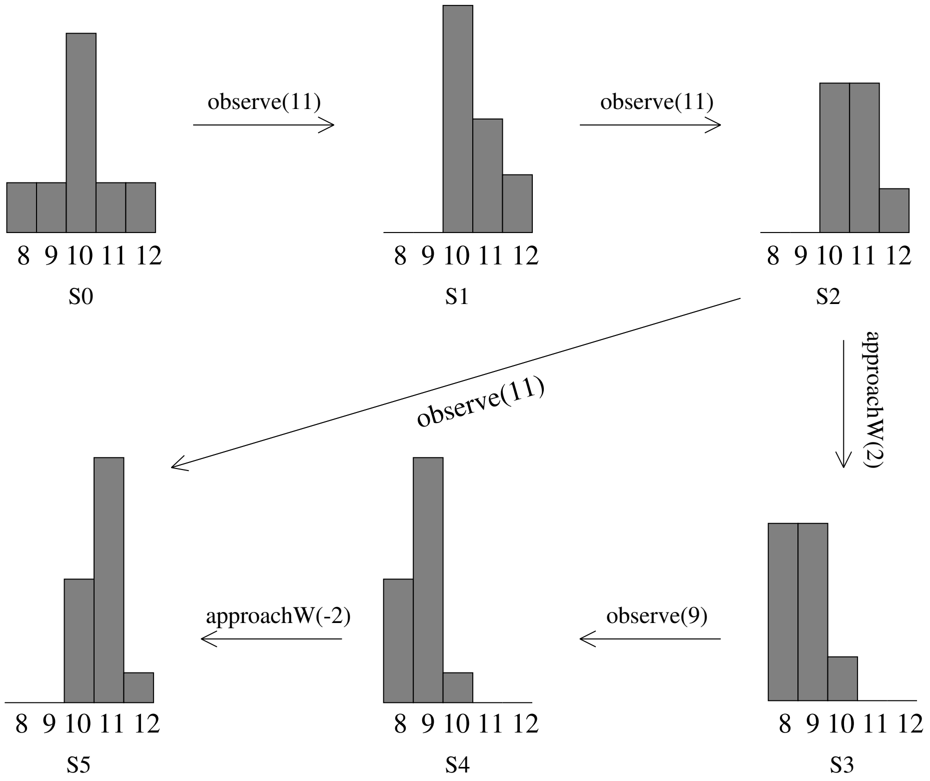

Figure 1: Example of Belief Update

correspond to sequences of sensing actions, and thus are handled by a simple iteration of Eq. 4. The independence of a sensor reading from all of the previous readings is implied by assumption I and by the fact that the sensors do not change the value being observed (this is captured in the successor-state axiom for the sensed fluent). As a result, after a sequence of sensing actions, the agent will come to have greater or less certainty about the value of the sensed fluent, dependent on whether or not the sequence of readings agree or not.

Example 5.2: Suppose that the agent executes another read action in the state S_1. Further, suppose that the agent observes the same value as before 11, and let $S_2 = do(\mathsf{observe}(11), S_1)$. Then, another application of Eq. 4 (applied to the agent's beliefs in S_1), yields the belief distribution shown in Figure 1. That is, the agent's beliefs have converged more tightly around the value 11, since it has now sensed that value twice.

As mentioned briefly in the previous section, the manner in which probability mass is transferred when ordinary actions like approachW also yields appropriate changes to the agent's beliefs.

Example 5.3: Suppose that the agent is in state S_2, and then moves 2 units closer to the wall. Let $S_3 = do(\mathsf{approachW}(2), S_2)$. Then, the successor state axiom for *p* and *wallDist* imply that the agent's beliefs are shifted to worlds in which it is 2 units closer to the wall. Hence, for all y, $\text{BEL}(\mathit{wallDist} = y - 2, S_3) = \text{BEL}(\mathit{wallDist} = y, S_2)$. The agent's shifted beliefs are shown in the diagram. This is exactly how one would expect the agent's beliefs to change after moving closer to the wall.

Furthermore, changes in the agent's beliefs due to ordinary actions integrate correctly with sensing actions.

Example 5.4: Suppose that the agent again executes a read action in S_3, and observes the value 9. Let $S_4 = do(\mathsf{observe}(9), S_3)$. This reading is consistent with its previous readings of 11 since the agent has moved 2 units closer to the wall. Hence, as shown in the diagram, it results in a further tightening of the agent's beliefs, around the value 9. If the agent subsequently moves back from the wall by 2 units, executing an $\mathsf{approachW}(-2)$ action ($S_5 = do(\mathsf{approachW}(-2), S_4)$), its beliefs will then be clustered around 11, as shown on the diagram.

Intuitively, since the agent's approachW action incurs no error, we would expect that if the agent had sensed the value 11 in situation S_2 without moving forward and then backward an equal amount, then its beliefs about the distance to the wall should be identical. Our model respects this intuition, as indicated in the diagram by the diagonal arrow from S_2 to S_5.

Finally, we can observe that if the agent executes an action that has no effect on a particular fluent, then that action will cause no change in the agent's beliefs about that fluent. For example, if the agent executes a drop action that has no effect on its distance to the wall, it will have exactly the same beliefs about the distance to the wall in the successor state. This again arises from the direct transfer of probability mass to the successor states, all of which have exactly the same distance to the wall as before.

6 Conclusion

We have demonstrated that noisy perception can be modeled in the situation calculus by a simple extension of previous work. Although the resulting formalism is limited in some ways, e.g., noisy effectors would require a non-trivial extension to our current approach, it does succeed in providing an interesting integration of noisy perception and ordinary actions. Most importantly, it succeeds in capturing some key features of this interaction that any more extensive formalism would also have to capture.

Much of our approach can be exported to alternate formalisms. For example, instead of the situation calculus a modal logic could have been used. Similarly, the probabilistic component could be replaced with an alternate formalism, like Dempster-Schafer belief functions, or fuzzy measures. All that would be required is to replace the functional fluent p and axioms for BEL with fluents and axioms to support an alternate measure of belief. The likelihood functions could then be replaced with non-probabilistic functions to support an alternate rule of belief update.

As for future work, apart from addressing limitations of the formalism, there is its application in high-level agent control. In the GOLOG work mentioned in the introduction, the ability of an agent to execute a program depends on what it *knows* about the truth value of the test conditions in that program [LLLS95]. When an agent only has a degree of belief in the truth of a test condition in a program, it is much less clear what it ought to do. A suitable programming formalism in this case remains to be developed.

References

[Bac90] F. Bacchus. *Representing and Reasoning with Probabilistic Knowledge*. MIT Press, Cambridge, Mass., 1990.

[FHMV95] R. Fagin, J. Y. Halpern, Y. Moses, and M. Y. Vardi. *Reasoning about Knowledge*. MIT Press, Cambridge, Mass., to appear, 1995.

[Haa87] A. R. Haas. The case for domain-specific frame axioms. In F. M. Brown, editor, *The Frame Problem in Artificial Intelligence: Proceedings of the 1987 workshop*, pages 343–348. Morgan Kaufmann, San Mateo, California, 1987.

[Hal90] J. Y. Halpern. An analysis of first-order logics of probability. *Artificial Intelligence*, 46:311–350, 1990.

[Har79] D. Harel. *First-Order Dynamic Logic*. Lecture Notes in Computer Science, Vol. 68. Springer-Verlag, Berlin/New York, 1979.

[HT93] J. Y. Halpern and M. R. Tuttle. Knowledge, probability, and adversaries. *Journal of the ACM*, 40(4):917–962, 1993.

[Lew76] D. Lewis. Probability of conditionals and conditional probabilities. *Philosophical Review*, 83(5):297–315, 1976.

[LLLS95] Yves Lespérance, Hector J. Levesque, Fangzhen Lin, and Richard B. Scherl. Ability and knowing how in the situation calculus. Submitted to IJCAI-95, 1995.

[LLR95] Hector J. Levesque, Fangzhen Lin, and Ray Reiter. Defining complex actions in the situation calculus. Dept. of Computer Science, Univ. of Toronto, in preparation, 1995.

[LR94] Fangzhen Lin and Raymond Reiter. State constraints revisited. *Journal of Logic and Computation*, 4(5):655–678, 1994.

[MH69] John McCarthy and Patrick J. Hayes. Some philosophical problems from the standpoint of artificial intelligence. In *Machine Intelligence 4*, pages 463–502. Edinburgh University Press, 1969.

[Moo85] R. C. Moore. A formal theory of knowledge and action. In J. R. Hobbs and R. C. Moore, editors, *Formal theories of the common sense world*, pages 319–358. Ablex Publishing, Norwood, NJ, 1985.

[Pea88] Judea Pearl. *Probabilistic Reasoning in Intelligent Systems*. Morgan Kaufmann, San Mateo, California, 1988.

[Ped89] E. P. D. Pednault. ADL: Exploring the middle ground between Strips and the situation calculus. In R. J. Brachman, H. J. Levesque, and R. Reiter, editors, *Proceedings of the International Conference on Principles of Knowledge Representation and Reasoning*, pages 324–332, 1989.

[Pin94] J. A. Pinto. *Temporal Reasoning in the Situation Calculus*. PhD thesis, Department of Computer Science, University of Toronto, 1994.

[Rei91] Raymond Reiter. The frame problem in the situation calculus: A simple solution (sometimes) and a completeness result for goal regression. In Vladimir Lifschitz, editor, *Artificial Intelligence and Mathematical Theory of Computation*, pages 359–380. Academic Press, 1991.

[SL93] Richard B. Scherl and Hector J. Levesque. The frame problem and knowledge-producing actions. In *Proceedings of the AAAI National Conference*, pages 689–695, 1993.

[Sch90] Schubert L. K. A monotonic solution to the frame problem in the situation calculus: An efficient method for worlds with fully specified actions. In H. E. Kyburg, R. P. Loui, and G. N. Carlson, editors, *Knowledge Representation and Defeasible Reasoning*, pages 23–67. Kluwer Academic Press, London, 1990.