Knowledge, Belief, and Noisy Sensing in the Situation Calculus

by

Patricio D. Simari

A thesis submitted in conformity with the requirements
for the degree of Master of Science
Graduate Department of Computer Science
University of Toronto

Abstract

Knowledge, Belief, and Noisy Sensing
in the Situation Calculus

Patricio D. Simari
Master of Science
Graduate Department of Computer Science
University of Toronto
2004

The extension to the situation calculus presented by Bacchus *et al.* formalizes the concept of noisy actions and shows how an agent can update its beliefs, which are modeled probabilistically, when relying on noisy sensors and effectors. The extensions of Scherl and Levesque and Shapiro *et al.* also model knowledge and belief. While assuming noiseless actions and dealing with boolean beliefs, these frameworks support properties of knowledge and belief such as introspection about current and past beliefs. Here, it is shown how such properties of belief can be formalized and supported in the probabilistic Bacchus *et al.* extension. In addition, the concept of sensor coarseness is introduced and it is shown how it can be modeled in the Bacchus *et al.* framework. Finally, it is shown that the Bacchus *et al.* framework can function in a way which is equivalent to using conditional probability densities to combine noisy sensor readings.

Acknowledgements

I would like to thank my supervisor Hector Levesque for all his support, ideas, and many helpful meetings and discussions. I am also grateful to Sam Roweis for reading this work and offering many corrections and suggestions.

I would also like to thank Diego Macrini for his friendship and advice, both technical and non-technical, and Laura Casole, for her love and encouragement.

Finally, I would like to thank my father Guillermo, my mother Amalia, and my brother Gerardo for all of their caring, love and support.

Contents

Chapter 1

Introduction

Russell and Norvig, in [18], define an agent as anything that can be viewed as perceiving its environment through sensors and acting upon that environment through actuators. Poole, in [15] also adds to this definition that, at any time, the agent has prior knowledge about the world it is in, past experience that it can learn from, goals that it must try to achieve or values about what is important, and observations about the current environment and itself.

Different areas of AI deal with different aspects pertaining to intelligent agents. For example, planning theory studies how agents can, based on their model of the world, obtain sequences of actions that achieve a certain goal. Decision theory studies how agents can make optimal decisions given what they sense and their current model of the world, and so on. This work, deals mainly with the aspect of agent beliefs: informally, those things that the agent considers to be true about its world

In practical situations, an agent will need to interact with a world that is not completely known to it. It will need to expand its internal model of the world with which it interacts by sensing. For example, an agent may require information regarding its distance to the nearest wall, or other obstacle, for which it may use a sensor, such as a sonar device.

However, such sensors are never completely accurate. If an agent receives a value indicating a distance of 0.501 meters to the nearest wall, this need not be the true value. Upon using the sensor a second time the agent may read a value of 0.502 or 0.499 meters. This is due to the fact that sensors such as the one described are subject to noise.

Another problem that may arise for the agent is that the conditions of the world it inhabits may change as a consequence of factors other than its own actions. For instance, a robot could collide with another object and lose track of its position, or even be moved by another agent, possibly a person. Consider, for example, the Robocup scenario, in which a robot player can be removed from the field, or the case of a space probe which is hit by debris. In scenarios of this type a considerable un-modeled positioning error is likely to occur. The problem of dealing with scenarios such as these is known as *the kidnapped robot problem.* [1] Frameworks that attempt to tackle this problem must give the agent the ability to unlearn its now incorrect beliefs and to determine its new position, preferably quickly. Each framework must define (as we will see in subsequent chapters) what it means for an agent to hold certain beliefs, and how these beliefs can change as a result of sensing.

The extension to the situation calculus (a well studied theory of action which we will review in chapter 2) presented by Bacchus *et al.* in [1] addresses these issues. It formalizes the concepts of noisy actions and shows how the agent can update its beliefs, which are modeled probabilistically, when relying on noisy sensors and effectors.

Other frameworks also address the formalization of knowledge and belief within the situation calculus. For example, the framework of Scherl and Levesque presented in [19], on which the work of Bacchus *et al.* builds, formalizes knowledge [2] and knowledge

[1]The notion of the kidnapped robot was introduced by Sean P. Engelson and Drew V. McDermot in [4]. An in-depth exploration of the various specific approaches put forward to deal with this problem is beyond the scope of this thesis. The interested reader is referred to [5], [9], [21], and [22].

[2]Knowledge in this framework is analogous to the concept of belief in the classic possible worlds model. However, Scherl and Levesque speak of *knowledge* to indicate full certainty in the belief. This full certainty in beliefs is implied in the possible worlds model, but as we will see later on it need not always be the case.

producing actions. It is based on the possible-worlds approach to belief, in which an agent has a set of worlds that it considers possible. Given a particular statement, it may hold in some worlds while in others it may not. An agent is said to *know* the statements that hold in all worlds the agent considers possible. The work focuses mainly on the theoretical properties of knowledge and assumes noiseless actions.

The framework of Shapiro *et al.* presented in [20] introduces belief to the situation calculus, while still assuming noise free sensing. Within this framework, in contrast to the one of Scherl and Levesque, an agent may have mistaken beliefs. Its contributions are the situation calculus formalization of belief and theoretical results showing how the framework allows for belief update (change in beliefs as a result of perceived changes in the world), belief revision (changes in beliefs without changes to the world), belief introspection (the agent knows what its beliefs are) and awareness of mistaken beliefs (the agent will realize that it believed something which was false.)

In both the framework of Scherl and Levesque as well as that of Shapiro *et al.*, knowledge and belief (respectively) are boolean in nature, rather than probabilistic. As one of the main contributions of this work we seek to show that the framework of Bacchus *et al.*, while dealing with probabilistic beliefs, can also model concepts such as belief introspection, previously held beliefs, current beliefs about the past, awareness of belief change and belief in the fact that previously held beliefs were mistaken. Moreover, we prove that in this framework, under certain conditions, the agent will have full introspection of its beliefs. This is to say, if an agent believes a certain statement with degree .73, then it will believe with degree 1 that it believes said statement with degree .73. Similar results are shown for the introspection of previously held beliefs, current beliefs about the past and belief change.

While it is true that the Bacchus *et al.* framework was created in order to model noisy sensing and how belief changes as a result of it, it can also model sensors which are noise free. We prove that under such conditions, the framework is well behaved with regards to belief, modeling the agent as having full certainty of the sensed value. For example, if an agent uses a sensor modeled as noiseless to ascertain its distance to the nearest wall and receives a value of 0.501 meters, then it will believe with degree 1 that its distance to the wall is 0.501 meters.

Another problem addressed in this work is that of sensor coarseness. In practical scenarios, it is possible for an agent to be equipped with a sensor that returns values which are "coarser" than the agent's beliefs. For example, when sensing its distance to the nearest wall, the agent may receive values in centimeters. Meanwhile, the agent may be able to have beliefs about distance in millimeters. We will show how the framework of Bacchus *et al.* can be adapted to handle such a scenario.

Finally, we compare the framework of Bacchus *et al.* to using conditional probability densities to combine noisy sensor data and show that the Bacchus *et al.* framework can function in a way which is equivalent.

The rest of this work is organized as follows: in chapter 2 we discuss the necessary background which we divide into two main sections. In the first section, we discuss the origins of the possible-worlds approach to belief and then proceed to give both semantic and proof theoretic characterizations. We also discuss variations of the main formalization and finish the section by exposing some issues that arise when combining belief operators with quantification. In the second section, we give an overview, firstly, of the basic situation calculus and then, of the three main publications on which this work is based, namely, the situation calculus extensions of Scherl and Levesque for representing knowledge ([19]), of Shapiro *et al.* for handling belief change ([20]), and of Bacchus *et al.* for handling uncertainty ([1]).

In chapter 3 we show how properties of belief can be formalized for the probabilistic framework of Bacchus *et al.*. These include belief revision under the assumption of accurate sensing, introspection of current beliefs, introspection of past beliefs, awareness of belief change, and belief in mistaken beliefs. In this section, we also include proofs showing that the agent has full introspection of its beliefs, both present and past, as well as its degree of belief change. Finally, we give an example in which each of the given properties is illustrated.

Chapter 4 introduces the concept of sensor coarseness. We show how it can be modeled in the framework of Bacchus *et al.* and discuss the consequences of coarseness for belief update. We illustrate the concept and conclude the chapter with an example.

In chapter 5 we give an example of using conditional probability densities to combine noisy sensor readings, and we show that the Bacchus *et al.* framework can function in a way which is equivalent. Here too we present illustrative examples.

Finally, chapter 6 presents the conclusions of our work.

Chapter 2

Background

2.1 The possible worlds approach to belief

In this section we give a brief overview of the origins of the possible worlds approach to belief. For a more in depth coverage of the topics discussed the reader is referred to texts such as [2], and [10], as well as [11], a survey which also provides a brief overview.

2.1.1 Basic notion

The basic idea behind the possible world's approach to belief is that there is a set of worlds that the agent considers possible. Given a certain proposition, it may be the case that it holds in some of these worlds, while in others it may not. An agent is then said to believe a given proposition if it holds in all worlds that it considers possible.

Jaakko Hintikka was one of the first to work on the logical formalization of the notions of belief and knowledge, providing a basis for the research that was to follow in this area. It was he who proposed the idea of using possible worlds semantics for said task.

Early work in this direction can be found in [6], including the introduction of the idea of an accessibility relation (which he called the alternativeness relation.) However it is in [7] that Hintikka formalizes his ideas into the now standard notion of possible worlds.

2.1.2 The logic used

In this overview, we will concentrate most of our focus on the propositional case, although toward the end of this section we will note some aspects that arise with quantification.

The logic used is simply propositional logic, with the addition of a modal operator B. In other words, the language of the logic contains:

- A set of propositional letters: p, q, r, $\ldots$
- Logical connectives: $\vee$ (disjunction), $\neg$ (negation)
- Modal operator: B
- Parentheses: (,)

Likewise, the formulas of the language are as in classic propositional logic, with the addition of the modal operator B as follows:

1. A propositional letter is a formula
2. If α and β are formulas, then so are $\neg\alpha$, $(\alpha \vee \beta)$, and $B\alpha$.

As is usual, when there is no ambiguity, parentheses may be omitted.

Other logical connectives, such as conjunction and material implication [1], $\wedge$ and $\supset$ respectively, can then be defined in terms of the primitive connectives.

- $(\alpha \wedge \beta) \stackrel{\text{def}}{=} \neg(\neg\alpha \vee \neg\beta)$
- $(\alpha \supset \beta) \stackrel{\text{def}}{=} (\neg\alpha \vee \beta)$

[1] Refers to statements of the form "if ..., then ...", or "implies"

2.1.3 Semantic characterization

Earlier, we mentioned that an agent will have a set of worlds he considers possible and it will be said to believe a proposition when said proposition holds in all possible worlds.

Saul Kripke introduces in [8] a semantic characterization of this notion of belief. In order to do so, Kripke defines a possible worlds model M as a triple of the form $\langle W, \pi, R \rangle$, where

- W is the set of all possible worlds.
- π is a function that maps proposition/world pairs to truth values. Given a proposition p and a world w, $w \in W$, $\pi(p, w) = t$ if p is true in w, and $\pi(p, w) = f$ if it is not.
- R is the accessibility relation. It is a binary relation over W defined so that, given $w, w' \in W$, wRw' if and only if in world w the agent considers w' to be a possible world. In other words, whenever wRw', as far as the agent knows while being in world w, he might be in world w'.

We may now define the support relation $\models$ between worlds and formulas as follows:

1. $w \models p$ iff $\pi(w, p) = t$, where p is a propositional letter.
2. $w \models \neg\alpha$ iff $w \not\models \alpha$.
3. $w \models \alpha \vee \beta$ iff $w \models \alpha$ or $w \models \beta$.
4. $w \models B\alpha$ iff for all w' such that wRw', $w' \models \alpha$

Finally, we may now define what it means for a formula to be satisfiable and valid in this context:

- A formula α is called satisfiable if there is some model $M = \langle W, \pi, R \rangle$ such that $w \models \alpha$ for some $w \in W$.

- A set of formulas $S = \{\alpha_1, \alpha_2, \ldots, \alpha_n\}$ is (simultaneously) satisfiable if there is some model $M = \langle W, \pi, R\rangle$ and some $w \in W$ such that $w \models \alpha_i$ for all i such that $1 \leq i \leq n$.

- A formula α is called valid (written $\models \alpha$), if for any model $M = \langle W, \pi, R\rangle$ and any $w \in W$, $w \models \alpha$.

2.1.4 Proof theoretic characterization

In the previous section we gave a semantic characterization of valid formulas. We will now give a proof theoretic characterization as presented in [11]. This formalization will constitute the basis of our discussion for the following subsection.

First, we state the axioms of the proof theory by describing them using axiom schemata:

- A1: All valid sentences in classical propositional logic.

- A2: $B(\alpha \supset \beta) \supset (B\alpha \supset B\beta)$

where α and β are well formed formulas of the propositional calculus. Next we state the rules of inference of the theory, i.e., we will characterize which formulas are theorems [2]:

- (*Modus Ponens*) From $\vdash \alpha$ and $\vdash (\alpha \supset \beta)$ infer $\vdash \beta$.

- (*Necessitation*) From $\vdash \alpha$ infer $\vdash B\alpha$.

If the accessibility relation R has no restrictions placed upon it, it can be shown that this proof theory is sound and complete with respect to the possible world semantics for belief. The logic resulting from this proof theory is called K.

[2] Hence the use of the symbol $\vdash$ instead of $\models$

2.1.5 Variations for the accessibility relation

K forms the basis of a group of logics (which include it) that capture different notions of knowledge and belief. We may arrive at the proof theories that result in these logics by placing certain restrictions on R and adding various other axioms to K. The following are three axioms that will be used to construct these proof theories:

- A3: $B\alpha \supset \alpha$
- A4: $B\alpha \supset BB\alpha$
- A5: $\neg B\alpha \supset B\neg B\alpha$
- D: $B\alpha \supset \neg B\neg\alpha$

The following table presents a summary of some of the more well known proof theories that arise from combining these theorems along with imposing restrictions on the accessibility relation, as well as the name of the resulting logic.

Proof theory	*Restrictions on R*	*Resulting logic*
K + A3	Reflexive	T
T + A4	Transitive, Reflexive	S4
S4 + A5	Reflexive, Transitive, Symmetric	S5
S4 - A3 + D	Transitive, Serial	Weak S4
S5 - A3 + D	Euclidean, Transitive, Serial	Weak S5

For the reader's convenience, we define each of type of relation:

- R is reflexive iff for all worlds u it is the case that uRu.

- R is transitive iff for any worlds u, v, and w, it is the case that if uRv and vRw then uRw.

- R is symmetric iff for any worlds u, v it is the case that uRv iff vRu.

- R is Euclidean iff for any worlds u, v, and w, it is the case that if uRv and uRw then vRw.

- R is serial iff for all worlds u there is some world v such that uRv.

2.1.6 Some issues regarding quantification

Previously in this section, we addressed the possible worlds approach to belief for the case of propositional logic. For a detailed exposition of the application of this approach to first order logic, the reader is referred to the sources cited at the beginning of this section. We will, however, briefly discuss some specific properties about belief with regards to quantification.

Let us first start with universal quantification in the case of the following two formulas involving belief:

1. $B(\forall x)\alpha \supset (\forall x)B\alpha$

2. $(\forall x)B\alpha \supset B(\forall x)\alpha$

Clearly, under any reasonable semantics, the first formula should be valid: if an agent believes "all ravens are black" to be true, then, given any particular raven x it should believe "x is black" to also be true.

Let us now consider the second formula, the converse of the first. This formula (known as the Barcan formula) is valid under semantics in which the domain is the same for every world and, consequently, *the agent is aware of what every element in the domain is.* However this need not be the case. In particular, the formula will not be valid in systems in which domains vary from world to world. Intuitively, an agent may believe, for every particular raven, that said raven is black. This does not mean, however, that the agent will believe "all ravens are black" to be true if it does not know that it has seen every raven.

We may also consider analogous formulas for the case of existential quantification:

1. $B(\exists x)\alpha \supset (\exists x)B\alpha$

2. $(\exists x)B\alpha \supset B(\exists x)\alpha$

In this case, it is the second formula that should be valid under any reasonable semantics. To continue using the example of ravens, if, for some x, the agent believes "x is a black raven", then the agent should also believe, "there is some black raven".

The first formula, however, is not valid. The agent may believe that "there is some black raven", without knowing specifically which raven it is. As another example, an agent may believe "someone in the group is a spy", without necessarily believing "x is a spy" for any particular person x.

2.2 Knowledge and action

2.2.1 The basic situation calculus: a brief overview

The situation calculus is a second order language designed for representing dynamically changing worlds. Here, we present a brief overview, describing those elements that shall be used in discussion during the remainder of this work. For an in-depth coverage, the interested reader is referred to [16]. It should also be stated that the examples presented in the following subsections are taken from said reference.

Situations and actions

Within the situation calculus, a possible state of the world is represented by a first-order term called a *situation*. A situation is either the initial situation, represented by the constant S_0, or the result of performing an *action* in a situation term. This performance is denoted by using the distinguished binary symbol *do* like so: $do(\alpha, s)$. Here, α denotes the action being carried out, while s is, recursively, a situation term.

Actions are denoted by function symbols accompanied by the appropriate parameters. For example, the action of putting object x on object y, could be represented as $put(x, y)$. The situation resulting from performing said action in a situation s, is denoted by $do(put(x, y), s)$.

Notice that a situation is, in essence, a history of the actions performed from the initial situation S_0. For example, $do(putDown(A), do(walk(L), do(pickup(A), S_0)))$ is a situation term denoting the sequence of actions $[pickup(A), walk(L), putDown(A)]$.

Fluents

In a *dynamic* world, it is natural that the values of functions will change from situation to situation. Analogously, while a relation may hold in one situation, it may or may not continue to hold in future situations.

Functions whose values vary from situation to situation are called *functional fluents*, and are denoted by function symbols taking a situation term as their last argument.

Relations whose truth values vary from situation to situation are called *relational fluents*. They are denoted by predicate symbols taking a situation term as their last argument.

In a world in which it is possible to paint objects, we might have a functional fluent $colour(x, s)$ denoting the colour of object x in situation s. Naturally the value of said fluent will depend on the colour of x at S_0 as well as the parameters of any *paint* actions carried out with x as the target.

In a mobile robot environment, there might be a relational fluent $closeTo(r, x, s)$, meaning that in situation s, the robot r will be close to the object x. Here too, whether this relation holds or not will necessarily depend on the position of r at S_0 as well as the parameters of any *move* actions carried out with r as the mover.

Action preconditions and effects

In order for an agent to be able to carry out a given action, it is necessary that certain conditions hold. These conditions will naturally depend on the specific action the agent is to execute and the situation said agent is in. For example:

1. A robot r can pick up an object x in situation s if, and only if, the robot is not holding any object, it is next to x, and x is not heavy.

2. A robot r can repair an object if, and only if, the object is broken, and said robot has glue.

These necessary conditions are *preconditions* for the corresponding actions. These actions can only be carried out if the preconditions hold.

In the situation calculus, preconditions for actions are modeled through the use of the predicate symbol $Poss$. Whether or not $Poss(a, s)$ holds is interpreted as whether or not it is possible to execute action a in situation s.

The following are the situation calculus formalizations of the corresponding English statements above[3]:

1. $Poss(pickup(r, x), s) \equiv [(\forall z)\neg holding(r, z, s)] \wedge \neg heavy(x) \wedge nextTo(r, x, s)$.

2. $Poss(repair(r, x), s) \equiv hasGlue(r, s) \wedge broken(x, s)$.

[3]Free variables are assumed quantified from the outside; i.e. the formula they are in is assumed to be contained in a larger formula in which there is a quantifier binding each free variable.

Just as there are special conditions that must hold before an action can be executed, there are also conditions that will hold after said action is executed. These conditions correspond to the *effect* of said action. For example:

1. If an object is fragile, then it will be broken when dropped.

2. When a robot repairs an object, the object is no longer broken.

3. Painting an object with a certain colour causes the object to be of that colour.

In the situation calculus, these effects are modeled as differences in the values of fluents in a given situation and the situation resulting from the execution of an action in said situation. Moreover, the effects are specified by what are called *effect axioms*.

The following are the situation calculus formalizations of the corresponding English statements above:

1. $fragile(x, s) \supset broken(x, do(drop(r, x), s))$.

2. $\neg broken(x, do(repair(r, x), s))$.

3. $colour(x, do(paint(x, c), s)) = c$.

The frame problem

The frame problem within the situation calculus was first posed by McCarthy and Hayes in [12] and then addressed by Raymond Reiter in [17]. We will briefly mention it here because the extensions of the situation calculus that will be discussed subsequently were devised in such a way as to inherit the solution to this problem presented by Reiter. For details, the interested reader is referred to the above mentioned publication, as well as [16].

In order to know which fluents are *not* affected by the performance of a given action, effect axioms are not sufficient. It is necessary to have what are known as frame axioms.

These specify the action invariants of the domain. Given that relatively few actions will affect the value of a given fluent, a very large number of these frame axioms will be needed. This is what is called the frame problem. Not only will the axiomatizer need to individually specify each and every one of these axioms, but also it will be difficult for the implementation to reason efficiently with such a large quantity of axioms.

In the situation calculus, the frame problem is solved by using three kinds of axioms:

1. Successor state axioms
2. Action precondition axioms
3. Unique name axioms for actions

Since later on we will be showing successor state axioms used in the frameworks we discuss, we take the opportunity to introduce their general form here.

Given a relational fluent F, the *successor state axiom for fluent F* is a formula of the form:

$$F(\vec{x}, do(a,s)) \equiv \gamma_F^+(\vec{x},a,s) \vee F(\vec{x},s) \wedge \neg\gamma_F^-(\vec{x},a,s)$$

where $\gamma_F^+(\vec{x},a,s)$ and $\gamma_F^-(\vec{x},a,s)$ are first-order formulas with free variables among $\vec{x},a,s$.

In the case of a functional fluent, the *successor state axiom for fluent F* is a formula of the form:

$$f(\vec{x}, do(a,s)) = y \equiv \gamma_f(\vec{x},y,a,s) \vee f(\vec{x},s) = y \wedge \neg\exists y' \gamma_f(\vec{x},y',a,s)$$

As an example, let us see the successor state axiom for fluent $Broken$:

$$\begin{aligned}
&Broken(x, do(a,s)) \equiv \\
&\quad \exists r\{a = drop(r,x) \wedge Fragile(x)\} \vee \\
&\quad \exists b\{a = explode(b) \wedge NextTo(b,x,s)\} \vee \\
&\quad Broken(x,s) \wedge \neg\exists r\{a = repair(r,x)\}
\end{aligned}$$

With the aforementioned three kinds of axioms, the frame axioms need not be specifically specified by the axiomatizer, and can be, instead, systematically generated in a parsimonious way and without omissions.

It should be noted that this solution does rely on two things: quantification over actions, and the assumption that relatively few actions affect a given fluent.

In the following subsections we describe extensions to the situation calculus that incorporate mechanisms for handling knowledge and belief. In each case, we will give a brief overview of said extensions, introducing any concepts or notation that will be required to comprehend later chapters of this work. The interested reader is referred to the references given in each subsection for full details of the described works.

2.2.2 Knowledge in the situation calculus

Building on the work of Moore [14], in [19], Scherl and Levesque incorporate the concept of knowledge and knowledge-producing actions into the situation calculus. This is done seamlessly by representing knowledge as a fluent. As such, it can be made to be affected by actions.

The knowledge fluent K is a binary relation over situations. This relation is understood to represent *accessibility* over possible worlds. In other words if $K(s', s)$ holds, then s' is said to be accessible from s, meaning that as far as the agent knows in situation s, he may actually be in situation s'.

Using the knowledge fluent, it is defined what it means for an agent to know something in a given situation, i.e., $Knows(\phi, s)$ for some formula ϕ and situation s.

Before giving the formal definition, let us introduce some notation. Since ϕ will normally contain fluents, the symbol *now* will be used as a placeholder for the situation argument of said fluents. By $\phi[s]$ we will denote the formula that results from substituting s for *now* in ϕ. The situation argument of fluents may be suppressed in the scope of a knowledge (or in subsequent sections, belief) operator, for readability.

In essence, it will be said that the agent *knows* ϕ in a situation s if ϕ holds in every situation s' accessible from s according to the fluent K. Formally

$$Knows(\phi, s) \stackrel{\text{def}}{=} \forall s'.K(s', s) \supset \phi[s'].$$

Note that the expression on the left of the equality is simply an abbreviation for the formula on the right.

We should mention here that this concept is analogous to the concept of belief in the possible worlds model presented at the beginning of this chapter. However, Scherl and Levesque name it *Knows* to indicate full certainty in the belief. This full certainty in beliefs was implied in the possible worlds model described earlier, but as we will see in the frameworks discussed later on in this chapter, it need not always be the case.

An important difference between this framework and that of the modal logics discussed at the beginning of this chapter is the ability of the agent to incorporate knowledge by means of sensing actions. In [19] this type of action is introduced into the situation calculus as follows.

Let us assume we have a knowledge-producing action sense-f which determines whether or not the fluent f is true or not. In the framework, this action would have an associated sensing result function SR. In the case of this boolean fluent the result would be "yes" if the fluent f is true and "no" otherwise. We may then axiomatize the sensing result as follows:

$$\begin{aligned} &SR(\mathsf{sense\text{-}f}, s) = r \equiv \\ &\quad (r = \text{“yes''} \wedge f(s)) \vee (r = \text{“no''} \wedge \neg f(s)) \end{aligned}$$

It is important to note that sensing result axioms are specified for all actions. If the action is not a sensing action, then the sensing result is some arbitrary value which is always the same, for example, "OK".

Next, let us show the successor-state axiom for the K fluent:

$$K(s'', do(a, s)) \equiv$$
$$(\exists s', s'' = do(a, s') \wedge K(s', s) \wedge Poss(a, s') \wedge SR(a, s) = SR(a, s'))$$

Intuitively, only worlds in which action a was possible and the same sensing result would be obtained remain K-related to the successor situation for s. Note that since non sensing actions always have the same sensing result, K related situations will remain K related after said actions are carried out.

In [19], Scherl and Levesque handle the knowledge fluent and knowledge-producing actions in such a way as to avoid the frame problem. This is done by proving that as a result of the specification, knowledge-producing actions do not affect fluents other than the knowledge fluent, and that actions that are not knowledge-producing only affect the knowledge fluent as appropriate. This solution is based on, and is in fact an extension of, Reiter's approach to the frame problem as presented in [17]

It is also shown in [19] how memory emerges as a consequence of the framework. Things known within a certain situation will remain known in successor situations, unless the opposite is specifically chosen.

A form of regression, examined by Reiter, is also shown to hold for knowledge-producing actions. By means of this regression, reasoning about knowledge and action in future situations is reduced to reasoning about knowledge in the initial situation. After this reduction, standard theorem proving techniques for modal logics may be used.

Finally, it is shown in [19] how desirable properties such as positive introspection will hold for all situations, if they hold at the initial situation.

2.2.3 Belief revision in the situation calculus

Shapiro *et al.*, in [20], present a way in which the situation calculus can be used to deal with belief, nested belief and belief update.

The work presented is an extension of the work of Scherl and Levesque in [19]. While Scherl and Levesque contemplate only belief expansion, the theory presented by Shapiro *et al.* allows for belief revision as well. What this means is that it is possible for an agent using this framework to be in a situation in which it believes a certain formula ϕ, acquires new information which causes it to then believe $\neg\phi$ and, after even more information is acquired, once again believe ϕ to be the case. Note that this was not the case in the framework of Scherl and Levesque discussed previously.

Given that within this framework beliefs change as a result of action, iterated belief change is naturally supported as being the result of a sequence of actions. Also the framework supports introspection: The agent can reason about what it believed in the past, what it believes in the present, and what it may believe in future situations resulting from present actions. Also, an agent will be able to tell the difference between an update and a revision. In the first case an agent will believe that it was correct about its past beliefs, while in the second case, the agent will believe that what it previously believed was wrong.

The essence of this framework is as follows. As in the case of Scherl and Levesque, an accessibility relation over situations is defined. In this framework, the operator is called B instead of K, given that it deals with belief instead of knowledge, but works in the same way. In addition, however, a plausibility relation over individual situations pl is defined. Given a situation s, $pl(s)$ indicates the plausibility of this situation according to the agent. It ranges over the natural numbers, lower numbers indicating a higher plausibility. This pl function is specified over initial situations and is then left unchanged for future situations. This is achieved through a simple successor state axiom.

Now, belief is defined in such a way that an agent will believe a proposition ϕ to hold in a situation s if ϕ holds in the most plausible B-related situations. Formally, this belief operator is defined as follows:

$$Bel_{Shap}(\phi, s) \stackrel{\text{def}}{=} \forall s'[B(s', s) \wedge (\forall s''.B(s'', s) \supset pl(s') \leq pl(s''))] \supset \phi[s'].$$

With this definition, the agent will always use only those accessible situations with the currently lowest *pl* value (i.e. the ones considered most plausible) in order to determine its beliefs. As soon as all situations with this value become inaccessible through the result of sensing and the application of the successor state axiom for B, a new subset of situations will be used for the purposes of determining belief: those with the next lowest *pl* value. In this way the truth value of Bel_{Shap} for a given ϕ may change from one situation to the next.

Shapiro *et al.* mention that Val and Shoham also present in [3] an extension of the situation calculus for belief revision that deals with both revision and update. However, their framework does not allow for nested belief, so introspection and mistaken belief cannot be dealt with.

We should mention that, within the Shapiro *et al.* framework, sensors are assumed to be completely reliable. This means that the agent will only revise its beliefs with true formulas.

Finally, [20] includes a comparison with AGM, KM and DP, showing that the theory satisfies most of the postulates of these frameworks.

2.2.4 Handling uncertainty

In [1], Bacchus *et al.* present a frame work which incorporates probabilistic reasoning into the situation calculus. The basis for doing this is that, in the real world, an agent's sensors and effectors are not completely reliable. The paper makes this clear by presenting an example: a robotic agent trying to determine the distance to the nearest wall may receive a sensor reading of 3.1 meters. However, because the sensor may introduce noise, the actual distance to the wall may be slightly greater or lesser. Similarly there may be inaccuracies in an agent's effectors. When the agent tries to move 0.7 meters toward the wall, the actual distance moved may be 0.8 due to these inaccuracies.

With the addition of probabilistic beliefs, the agent is able to reason about the values returned by noisy sensors and those produced by noisy effectors. In this extended framework, even though the reading obtained may not be completely accurate, the agent will be able to increase its belief in the fact that the actual value is the one returned by the sensor. Respectively, it will also be able to update its beliefs accordingly after carrying out actions which rely on noisy effectors.

The framework is built upon the variant of the situation calculus which supports the solution to the frame problem proposed by Reiter in [17] as well as the work of Scherl and Levesque in [19]. Based upon this previous work, probabilities are added and the effects of actions upon an agent's probabilistic beliefs are modeled.

The knowledge fluent K of Scherl and Levesque is present in this framework, but in addition, Bacchus *et al.* introduce a new fluent p. The goal of this introduction is to allow the agent to associate with each possible world a numeric value indicating the agent's subjective degree of belief that that is the actual world. So, $p(s', s)$ denotes the relative weight the agent places on s' being the actual world, when it is in situation s. We should mention that weights are expected to be non-negative, and that situations considered impossible are associated a weight of 0. This is insured in the framework by an initial constraint over situations p-related to the initial situations and by the successor-state axiom for p.

Now, given a formula ϕ, the agent's degree of belief in ϕ is calculated as the total weight of all worlds it considers possible where ϕ holds, normalized by the total weight of all worlds it considers possible.

Formally, this is defined as follows. For any formula ϕ and situation s

$$Bel(\phi, s) \stackrel{\text{def}}{=} \frac{\sum_{\{s':\phi[s']\}} p(s', s)}{\sum_{s'} p(s', s)}$$

An appendix of [1] contains a formalization of these summations expressing them in first order logic and using sorted variables.

In order to model noisy sensors and effectors, Bacchus *et al.* provide a way to formalize nondeterministic actions in the situation calculus. Nondeterministic actions are formalized as the execution of one of many possible deterministic actions, without knowing which one was actually carried out.

An example presented in [1] is that of a noisy action called $\mathsf{noisy\text{-}advance}$. Assuming a one-dimensional world, when an agent carries out the action $\mathsf{noisy\text{-}advance}(x)$, it will actually advance a distance y, approximately equal to x but whose exact value is unknown to the agent. In order to model the action $\mathsf{noisy\text{-}advance}$, a deterministic action of $\mathsf{advance}(x, y)$ is used, where x is the distance the agent is trying to move while y is the actual distance moved. Note that the agent may specify x, but has no control over y, so the agent cannot directly execute an action $\mathsf{advance}(x, y)$.

Assuming there is a fluent *position* whose value can only be changed by $\mathsf{advance}$, a successor state axiom for it can now be written thusly:

$$\begin{aligned}
&position(do(a, s)) = z \equiv \\
&\qquad \exists x, y\{a = \mathsf{advance}(x, y) \wedge z = position(s) + y\} \vee \\
&\qquad z = position(s) \wedge \neg\exists x, y\{a = \mathsf{advance}(x, y)\}
\end{aligned}$$

The nondeterministic action $\mathsf{noisy\text{-}advance}(x)$ can now be modeled as the execution of $\mathsf{advance}(x, y)$ with a nondeterministic choice of y. This is denoted in [1] as follows:

$$\mathsf{noisy\text{-}advance}(x) \stackrel{\text{def}}{=} \pi y.\mathsf{advance}(x, y).$$

Using action precondition axioms (involving the predicate symbol $Poss$), a relationship between x and y may be established so as to limit the amount of $\mathsf{advance}$ actions that may arise.

Although the agent may not know the exact effect of its actions, there is a correlation between the intended effect and the resulting one. This correlation is specified in the framework through the use of action-likelihood functions.

These are of the form

$$\ell(a, s) = z \equiv \phi(a, z, s).$$

The formula ϕ characterizes the conditions under which action a has likelihood z in situation s.

In the framework of Scherl and Levesque covered previously, when an agent performed a sensing action on a fluent it received the fluents true value. Likewise, when carrying out an action of, say, moving, the amount moved was modeled as always being the amount specified by the agent. In the Bacchus *et al.* framework, however, it is necessary to model the fact that the agent does not know the actual result of the noisy actions it performs. For this purpose the concept of observation indistinguishability is introduced.

Let us consider, for example, the case of the noisy action $\mathsf{advance}$. Any two actions of the form $\mathsf{advance}(x, y)$ and $\mathsf{advance}(x, y')$ will be observationally indistinguishable for the agent. All the agent knows is that the value y (or y') will be a value compatible with x according to the likelihood function for $\mathsf{advance}$. Formally, when two actions are observationally indistinguishable is determined by observation indistinguishability axioms. These are defined through the use of the predicate symbol OI, and will have the following form:

$$OI(a, a', s) \equiv \phi(a, a', s).$$

There will be one for each action a, and in each case a formula ϕ will characterize the relationship between a and a' that makes them indistinguishable to the agent.

Through the use of precondition axioms, action-likelihood functions, observation indistinguishability, and the knowledge fluent K, Bacchus *et al.* specify the successor state axiom for the fluent p.

$$\begin{aligned}
&p(s'', do(a, s)) = z \equiv \\
&\qquad \exists a', s'\{s'' = do(a', s') \wedge Poss(a', s') \wedge OI(a, a', s) \wedge K(s', s) \wedge z = p(s', s) \times \ell(a', s')\} \\
&\vee \quad z = 0 \wedge \neg\exists a', s'\{s'' = do(a', s') \wedge Poss(a', s') \wedge OI(a, a', s) \wedge K(s', s)\}
\end{aligned}$$

In this way, given that belief is computed from the agent's p values, it is defined how belief changes as a result of carrying out nondeterministic actions.

With all of these elements, an agent operating under this framework will see its degree of belief in the value of a fluent change from situation to situation, and, it will change consistently with what is knows about the characteristics of the noise of the sensor or effector being used. What's more, through repeated sensings, the agent's certainty about the fluent's value will increase despite sensor noise. An example at the end of this subsection will illustrates this.

An important element included in the publication is that, if the following two conditions hold in the agent's formalization:

1. Sensing actions do not effect the world

2. The result of a sensing action depends only on the true value of the fluent being sensed and sensor noise, and is otherwise independent of the situation.

then the framework carries out belief update in a manner identical to standard Bayesian conditioning. For the proof of this fact, the reader is referred to the publication ([1]). We will however, include here the characterization of the agent's updated beliefs after reading a sensor value.

Let us assume we have a noisy sensing action noisy-sense-f to sense the functional fluent f, defined as:

$$\mathsf{noisy\text{-}sense\text{-}f} \stackrel{\text{def}}{=} \pi x, y.\mathsf{sense\text{-}f}(x, y)$$

Here, x is the value appearing on the agent's sensor and y is the actual value. The precondition axiom for the primitive action sense-f is:

$$Poss(\mathsf{sense\text{-}f}(x, y), s) \equiv y = f(s)$$

And, we assume that the action likelihood function ℓ is dependent only on the value of f and the value sensed:

$$\ell(\mathsf{sense\text{-}f}(x, y), s) = \ell(\mathsf{sense\text{-}f}(x, y))$$

Under these conditions, then in the situation $s^+ = do(\mathsf{sense\text{-}f}(x, f(s)), s)$ it Bacchus *et al.* show that it will hold that

$$Bel(f(now) = t, s^+) = \frac{Bel(f(now) = t, s)\ell(\mathsf{sense\text{-}f}(x, t))}{\sum_{t'} Bel(f(now) = t', s)\ell(\mathsf{sense\text{-}f}(x, t'))}$$

Example

The following example is taken, slightly modified, from [1].

Suppose that an agent is sensing its distance in meters to the nearest wall *distance* using a noisy-sense-distance action defined as

$$\mathsf{noisy\text{-}sense\text{-}distance} \stackrel{\text{def}}{=} \pi x, y.\mathsf{sense\text{-}distance}(x, y)$$

For the purposes of this example, let the action-likelihood axiom for sense-distance be

$$\begin{aligned}
&\ell(\mathsf{sense\text{-}distance}(x, y), s) = \\
&\qquad \textbf{if } x = y \textbf{ then } 0.5 \\
&\qquad \textbf{else if } |x - y| = 1 \textbf{ then } 0.25 \\
&\qquad \textbf{else } 0
\end{aligned}$$

Here, we are assuming that *distance* and the arguments to sense-distance sense distance can only take on integer values (i.e., this is the precision of these numbers). The axiom specifies that there is zero probability that the sensor will read a *distance* that is greater than 1 unit away from the true *distance*, and that the sensor noise is independent of other features of the situation.

Let the agent's initial beliefs regarding *distance* be given by:

- $Bel(distance(now) = t, S_0) = 1/8$ for $t \in \{8, 9, 11, 12\}$
- $Bel(distance(now) = 10, S_0) = 1/2$

Initially, the agent does not ascribe positive probability to any other possible value for *distance*. This distribution of beliefs for various values of *distance* in S_0 is shown on the top left-hand corner of figure 2.1.

Suppose that the agent senses its distance and observes a value of 11. If $S_1 = do(\textsf{sense-distance}(11, f(S_0)), S_0)$, a simple calculation shows that the agent's new beliefs in this situation will be:

- $Bel(distance(now) = 10, S_1) = 4/7$
- $Bel(distance(now) = 11, S_1) = 2/7$
- $Bel(distance(now) = 12, S_1) = 1/7$

This new distribution is shown on the top right-hand corner of figure 2.1. Since, with probability 1, the sensor returns a value that is within 1 unit of the true value, the agent now has a degree zero in the values 8 and 9.

Note in figure 2.1 how in situation S_1 the agent still believes that the most likely value of *distance* is 10, even though its sensor returned the value 11. This is a consequence of the agent's high prior belief in the value being 10.

The rest of figure 2.1 shows how the agent's beliefs evolve after performing two more sensings, both of which return the value 11. Note how, despite sensor noise, the agent's certainty in the value 11 increases through repeated sensing.

Some remarks on notation

In the rest of this work, we will often describe scenarios in which an agent's beliefs for some formula ϕ are normally distributed with mean μ and deviation σ in some situation s. In such cases, we will indicate this by saying that it holds that:

$$\forall x. Bel(\phi, s) = Normal((x - \mu)/\sigma)$$

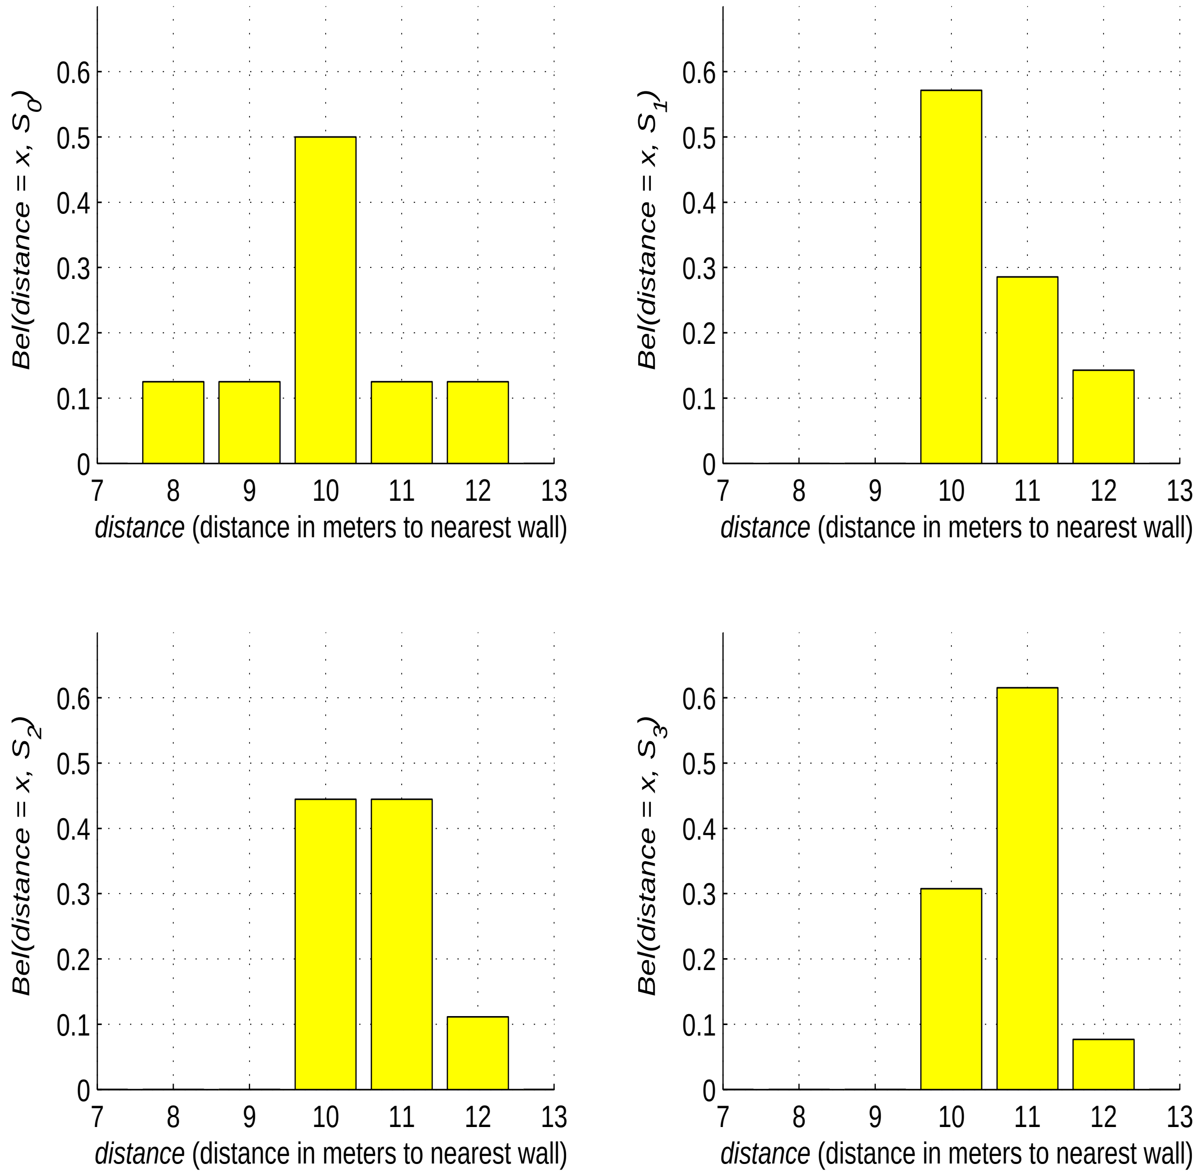

Figure 2.1: **Top left**: The agent's initial distribution of beliefs for various values of *distance* in situation S_0 **Top right**: Distribution for $Bel(distance = x, S_1)$, where $S_1 = do(\textsf{sense-distance}(11, f(S_0)), S_0)$. **Bottom left**: Distribution for $Bel(distance = x, S_2)$, where $S_2 = do(\textsf{sense-distance}(11, f(S_1)), S_1)$. **Bottom right**: Distribution for $Bel(distance = x, S_3)$, where $S_3 = do(\textsf{sense-distance}(11, f(S_2)), S_2)$.

where x is a free variable in ϕ.

We are assuming that $Normal$ is a discrete distribution. Action parameters will be finite precision numbers, and the value of $Normal(x)$ is the result of integrating the normal density function over the range of precision of x. We will assume that the domain theory contains a set of equations expressing the required values of $Normal$ according to the fixed precision being used.

Also, sometimes we will wish to state that the agent's beliefs are distributed in a way that is proportional to the sum of a Gaussian plus a small uniform distribution in some situation s. In such cases, given a formula ϕ, mean μ, deviation σ and small real number ϵ we will state that the following holds:

$$\forall x.Bel(\phi, s) \propto Normal((x-\mu)/\sigma) + \epsilon$$

where x is free in ϕ.

The symbol $\propto$ is read as "is proportional to". We need to say that the belief distribution is proportional, and not equal, because since it is a probability distribution, it is necessary that it sum to 1. Note that the use of $\propto$ adds no expressive power, given that it can be represented in first order logic simply by stating that it holds that

$$\exists r \forall x.Bel(\phi, s) = (Normal((x-\mu)/\sigma) + \epsilon)/r$$

These definitions and notational conventions will also be used when describing action-likelihood axioms.

Chapter 3

Properties of the Bacchus *et al.* extension

In [20], Shapiro *et al.* show that their extension of the situation calculus has the power to represent belief revision, belief update, introspection, and awareness of mistakes.

In this section, we will present similar results for the Bacchus *et al.* framework. In essence, we will formulate properties that are analogous to those presented in [20] and we will show that they hold. Specifically, we will show that, under certain conditions, the framework of Bacchus *et al.* handles:

- Revision with accurate sensing
- Introspection of current beliefs
- Introspection of past beliefs
- Awareness of degree of belief change
- Belief in mistaken beliefs

We will describe each of these in turn in the following sections. Before we proceed however, it is important to note that the results we will discuss in the following sections are valid only if it holds (unless otherwise noted) that the p value of any given pair of situations is only dependent on the target situation. For this, it is necessary to make the following assumption:

Assumption 3.0.1 *For any initial situations s_0, s'_0, s''_0 it holds that $p(s_0, s'_0) = p(s_0, s''_0)$*

This assumption is not unreasonable. In order to assign the p values for initial situations, one can imagine an agent having, for each possible fluent (that is independent of the values of other fluents), a probability distribution over the possible values of said fluent. Then, the probability of an initial situation s_0, is merely the joint probability of all (independent) fluents taking the values that they take in s_0. Clearly, the p value of any pair of initial situations will depend only on the target situation.

With this assumption we ensure that p values depend only upon the target situation for all pairs of situations, not just initial ones, as shown by the following theorem.

Theorem 3.0.1 *In an action theory where assumption 3.0.1 holds it also holds that*

$$\forall s^+, s'^+, s''^+ . p(s^+, s'^+) = p(s^+, s''^+)$$

Proof Let us assume we are dealing with an action theory in which assumption 3.0.1 holds and let s^+, s'^+, s''^+ be any three p-related situations.

If any of the three is an initial situation, then, given that they are p-related, all three will be initial situations in which case it trivially holds that $p(s^+, s'^+) = p(s^+, s''^+)$ from assumption 3.0.1.

Let us assume that they are not initial situations and let us further assume that the property holds for all previous situations.

Since s^+, s'^+, s''^+ are not initial situations, then there exist situations s, s', s'' and actions a, a' and a'' such that $s^+ = do(a, s)$, $s'^+ = do(a', s')$, $s''^+ = (a'', s'')$.

By the successor state axiom for p, it holds that $p(s^+, s'^+) = p(s, s') \times \ell(a, s)$ and $p(s^+, s''^+) = p(s, s'') \times \ell(a, s)$. But since s, s' and s'' are p-related situations previous to s^+,

s'^{+}, s''^{+}, then by hypothesis it holds that $p(s, s'') = p(s, s')$. Replacing above, we obtain $p(s^{+}, s''^{+}) = p(s, s') \times \ell(a, s)$. So, by transitivity, it holds that $p(s^{+}, s'^{+}) = p(s^{+}, s''^{+})$. $\blacksquare$

3.1 Belief revision with accurate sensing

In both [19] and [20] it is shown that in these frameworks, in the situation resulting from performing a sensing action on a given fluent, the agent knows the truth value of said fluent. In both of these frameworks, however, the sensors are assumed to be noise-free, always returning the true value of the sensed fluent.

Even though the Bacchus *et al.* framework was created to deal with noisy sensors, it has the expressive power to represent a sensor that is accurate; *i.e.* free of noise. Here, we will prove that, under the assumption of accurate sensing, an analogous result holds.

The way this property can be stated for this framework is as follows:

Property 3.1.1 *Assuming an agent has a* noise free *sensor for some fluent f, and $f(now) = x$ for some value x holds in the current situation, then the agent's degree of belief in $f(now) = x$ will be* 1 *after performing a sensing action on f.*

Note This property does not depend on assumption 3.0.1.

We formalize this property in the following theorem.

Theorem 3.1.1 *Given a fluent f and a noisy sensing action* noisy-sense-f *for said fluent defined as*

$$\textsf{noisy-sense-f} \stackrel{\text{def}}{=} \pi x, y.\textsf{sense-f}(x, y)$$

where the action-likelihood axiom for sense-f *is of the form*

$$\ell(\textsf{sense-f}(x,y),s) = \begin{cases} 1 & \textit{if } x = y = f(s); \\ 0 & o.w. \end{cases}$$

then it holds that

$$f(s) = x \supset Bel(f(now) = x, do(\textsf{sense-f}(x, f(s)), s)) = 1$$

Proof Let us assume $f(s) = x$ holds. By definition of Bel we may write

$$Bel(f(now) = x, do(\textsf{noisy-sense-f}(x), s)) \stackrel{\text{def}}{=} \frac{\sum_{\{s':f(s')=x\}} p(s', do(\textsf{noisy-sense-f}(x), s))}{\sum_{s'} p(s', do(\textsf{noisy-sense-f}(x), s))} \tag{3.1}$$

Now, by using the successor state axiom for p, we may write

$$p(s', do(\textsf{sense-f}(x, f(s)), s)) = p(s'', s) * \ell(a, s'')$$

for some action a and situation s'' such that $s' = do(a, s'')$, being a observationally indistinguishable from $\textsf{noisy-sense-f}(x)$. Then, by the way $\textsf{noisy-sense-f}$ is defined above and by definition of observation-indistinguishability, it must be the case that $a = \textsf{sense-f}(x, f(s''))$. So now we may write

$$p(s', do(\textsf{sense-f}(x, f(s)), s)) = p(s'', s) * \ell(\textsf{sense-f}(x, f(s'')), s'')$$

Using the above equality, we may rewrite (3.1) as

$$\frac{\sum_{\{s':f(s')=x\}} p(s'', s) * \ell(\textsf{sense-f}(x, f(s'')), s'')}{\sum_{s'} p(s'', s) * \ell(\textsf{sense-f}(x, f(s'')), s'')} \tag{3.2}$$

We may decompose the denominator of this expression into the following sum:

$$\sum_{s'} p(s'', s) * \ell(\textsf{sense-f}(x, f(s'')), s'') =$$
$$\sum_{s':f(s')=x} p(s'', s) * \ell(\textsf{sense-f}(x, f(s'')), s'') +$$
$$\sum_{s':f(s')\neq x} p(s'', s) * \ell(\textsf{sense-f}(x, f(s'')), s'')$$

Let us prove that

$$\sum_{s':f(s')\neq x} p(s'', s) * \ell(\textsf{sense-f}(x, f(s'')), s'') = 0$$

In order to do this, it suffices to show that the following holds

$$f(s') \neq x \supset p(s'', s) * \ell(\textsf{sense-f}(x, f(s'')), s'') = 0$$

In effect, let us assume $f(s') \neq x$ holds. By hypothesis, we also have that $f(s) = x$ holds. This implies that $f(s') \neq f(s)$.

But also, since a is a sensing action we have that $f(s') = f(do(a, s'')) = f(s'')$. So $f(s'') \neq f(s)$, and therefore $f(s'') \neq x$.

Hence, by definition of ℓ we know that

$$\ell(\textsf{sense-f}(x, f(s'')), s'') = 0$$

Therefore, $p(s'', s) * \ell(\textsf{sense-f}(x, f(s'')), s'') = 0$ for all situations s' such that $f(s') \neq x$ holds, and so

$$\sum_{s':f(s')\neq x} p(s'', s) * \ell(\textsf{sense-f}(x, f(s'')), s'') = 0$$

Using this result in our previous decomposition we now have

$$\sum_{s'} p(s'', s) * \ell(\textsf{sense-f}(x, f(s'')), s'') = \sum_{s':f(s')=x} p(s'', s) * \ell(\textsf{sense-f}(x, f(s'')), s'')$$

and so the numerator and denominator in expression 3.2 are equal. Therefore,

$$Bel(f(now) = x, do(\textsf{sense-f}(x, f(s)), s)) = 1$$

$\blacksquare$

3.2 Introspection of current beliefs

An important property of the extensions examined both in [19] and [20] is that of *belief introspection*. In the case of the Bacchus *et al.* framework, we will show a similar property holds. We should note that, given that beliefs are numeric in nature, it does not make sense to speak of negative introspection; the agent will simply have beliefs about its degree of belief in given formulas.

Property 3.2.1 *An agent is able to introspect its current beliefs with full certainty.*

The following theorem formalizes this property.

Theorem 3.2.1 *For all formulas ϕ and situations s, it holds that*

$$Bel(\phi, s) = x \supset Bel(Bel(\phi, now) = x, s) = 1$$

Proof For any formula ψ and situation s we know that $Bel(\psi, s)$ is an abbreviation for

$$\frac{\sum_{\{s':\psi[s']\}} p(s', s)}{\sum_{s'} p(s', s)}$$

So, in order for it to be true that $Bel(Bel(\phi, now) = x, s) = 1$, it must be the case that the numerator and the denominator of the previous expression evaluate to equal results, and thus $\psi[s']$ must hold in all situations s' such that $p(s', s) > 0$ is true.

Replacing ψ by $Bel(\phi, now) = x$, we must then show that, given a situation s', if $p(s', s) > 0$ holds then $Bel(\phi, s') = x$ holds.

Let us now consider any situation s' such that $p(s', s) > 0$ is true and show that it must also be true that $Bel(\phi, s') = x$.

In effect, we can rewrite $Bel(\phi, s) = x$ as

$$\frac{\sum_{\{t:\phi[t]\}} p(t, s)}{\sum_t p(t, s)} = x$$

But, from assumption 3.0.1 and theorem 3.0.1 we know that $\forall t.p(t, s) = p(t, s')$ is true, so replacing $p(t, s)$ for $p(t, s')$ in the previous equation, we derive that it holds that

$$\frac{\sum_{\{t:\phi[t]\}} p(t, s')}{\sum_t p(t, s')} = x$$

And finally, this can be abbreviated as $Bel(\phi, s') = x$ being true, which is what we wanted to show. ∎

3.3 Introspection of past beliefs

Let us introduce the notation $PrevBel(\phi, s) = x$ to denote that the agent believed ϕ with degree x in the situation immediately before s. Formally:

Definition

$$PrevBel(\phi, s) = x \stackrel{\text{def}}{=} \exists a, s'.s = do(a, s') \wedge Bel(\phi, s') = x$$

Property 3.3.1 *Agent is able to introspect the beliefs of the previous situation with full certainty.*

The following theorem formalizes this property.

Theorem 3.3.1 *For all formulas ϕ and situations s, it holds that*

$$PrevBel(\phi, s) = x \supset Bel(PrevBel(\phi, now) = x, s) = 1$$

Proof Let us assume we have a situation s, a formula ϕ, and a real number x such that $PrevBel(\phi, s) = x$ holds.

Analogously to the proof of theorem 3.2.1, in order to prove that $Bel(PrevBel(\phi, now) = x, s) = 1$ holds, it is sufficient to show that for any situation s' such that $p(s', s) > 0$ is the case, $PrevBel(\phi, s') = x$ holds.

Replacing $PrevBel(\phi, s) = x$ by its definition we then know by hypothesis that it is true that $\exists a, t.s = do(a, t) \wedge Bel(\phi, t) = x$.

Now, let us consider a situation s' for which $p(s', s) > 0$ is the case. We know that s is not an initial situation, and since s' is p-related to s we then know s' is not an initial situation. Therefore, it must be the case that there is an action a' and a situation t' such that $s' = do(a', t')$.

From the successor-state axiom for p, we know that $p(s', s) = p(t', t) \times \ell(a', t')$, and since we know that $p(s', s) > 0$ then it must also be the case that $p(t', t) > 0$. From this and the fact that $Bel(\phi, t) = x$, using reasoning analogous to the one found in the proof for theorem 3.2.1, we may conclude that $Bel(\phi, t') = x$ holds.

Thus, we have shown that $\exists a', t'.s' = do(a', t') \wedge Bel(\phi, t') = x$ holds. Hence, by definition of $PrevBel$, $PrevBel(\phi, s') = x$ holds, which is what we wanted to show. ∎

3.4 Awareness of belief change

In [20], Shapiro *et al.* show that in their framework, an agent is aware of mistakes it made in its beliefs. In other words, an agent may believe that, in a previous situation, it believed a certain formula ϕ held when, actually, it was the case that $\neg\phi$ held.

Given that we are dealing with noisy sensors, an agent will not know if it was actually mistaken in its past beliefs. However it can, for example, realize that its current sensor readings are returning values to which it had previously assigned a low degree of belief. Thus, although an agent may not be aware of a mistake *per se*, it will be aware that its beliefs regarding a certain formula have changed in some degree from the previous situation to the current one.

To this effect, we introduce the notation $BelChange(\phi, s) = \delta$ to denote that the agent's degree of belief in ϕ has changed by δ from the situation immediately previous to s to s itself.

Definition

$$BelChange(\phi, s) = \delta \stackrel{\text{def}}{=}$$
$$(\exists x_1 . PrevBel(\phi, s) = x_1) \wedge (\exists x_2 . Bel(\phi, s) = x_2) \wedge (\delta = x_2 - x_1)$$

Let us note that, since δ is the difference of two degrees of belief which vary between 0 and 1, δ itself will vary between -1 and 1. If δ is positive this means that the agent's degree of belief in ϕ has increased and, conversely, if δ is negative, this means that the agent's degree of belief in ϕ has dropped. The magnitude of δ tells us how radical the belief change is. A magnitude close to 0 indicates that the agent's degree of belief in ϕ remains similar, where as a magnitude close to 1 indicates that the agent's degree of belief regarding ϕ has changed radically.

Property 3.4.1 *An agent is able to introspect its degree of belief change with full certainty.*

The following corollary formalizes this property.

Corollary 3.4.1 *For all formulas ϕ and situations s, it holds that*

$$BelChange(\phi, s) = \delta \supset Bel(BelChange(\phi, now) = \delta, s) = 1$$

Proof Immediate from previous two theorems. *BelChange* is an abbreviation for a conjunction. Given that the previous theorems prove that, given the hypotheses, the agent's degree of belief in each of the conjuncts will be 1, then we may conclude that the agent's degree of belief in the conjunction will be 1.

3.5 Belief in mistaken beliefs

In the previous section we formalized how an agent may know about its distribution of beliefs in the previous situation. This corresponds to an agent knowing what it used to believe.

Now, we formalize a slightly different concept: how an agent may have beliefs in the present situation about what held in the past.

To this effect, we introduce the notation $BelThatInPrev(\phi, s) = x$ to denote that the agent now believes with some degree x that ϕ held in the previous situation.

Definition

$$BelThatInPrev(\phi, s) = x \stackrel{\text{def}}{=} \frac{\sum_{\{s' : \exists a, s''. s' = do(a, s'') \wedge \phi[s'']\}} p(s', s)}{\sum_{s'} p(s', s)}$$

Using this concept, we are now able to formalize the notion of an agent believing that its previous beliefs were wrong. This corresponds to an agent having a high degree of belief in the conjunction of the following two statements:

- It believed (with high confidence) in the previous situation that a certain formula ϕ was true.

- Said formula ϕ was false in that situation.

Using the *BelThatInPrev* notation we have just introduced, this scenario may be captured by stating that the following formula holds:

$$BelThatInPrev(Bel(\phi, now) \geq x_1 \wedge \neg\phi, s) \geq x_2$$

where s is the current situation and x_1 and x_2 represent reasonably high degrees of confidence. The example at the end of this chapter will help illustrate this concept.

Property 3.5.1 *An agent is able to introspect its current beliefs about what held in the previous situation with full certainty.*

The following theorem formalizes this property.

Theorem 3.5.1 *For all formulas ϕ and situations s, it holds that*

$$BelThatInPrev(\phi, s) = x \supset Bel(BelThatInPrev(\phi, now) = x, s) = 1$$

Proof Let us assume we have a situation s, a formula ϕ and a real number x for which $BelThatInPrev(\phi, s) = x$ holds.

Analogously to previous proofs, in order to show that $Bel(BelThatInPrev(\phi, s) = x) = 1$ holds, it is sufficient to show that for any situation s' such that $p(s', s) > 0$ is the case, then $BelThatInPrev(\phi, s') = x$ holds.

Replacing $BelThatInPrev(\phi, s) = x$ by its definition, we know by hypothesis that the following holds:

$$\frac{\sum_{\{t:\exists a,t'.t=do(a,t')\wedge\phi[t']\}} p(t, s)}{\sum_t p(t, s)} = x$$

Now let us consider a situation s' for which $p(s', s) > 0$ is the case. Given that s and s' are p-related, we know that the set of situations p-related to s is equal to the set of situations p-related to s'. Furthermore, by assumption 3.0.1 and theorem 3.0.1, we know that for all situations t it is the case that $p(t, s) = p(t, s')$. Therefore it must hold that

$$\frac{\sum_{\{t:\exists a,t'.t=do(a,t')\wedge\phi[t']\}} p(t, s)}{\sum_t p(t, s)} = \frac{\sum_{\{t:\exists a,t'.t=do(a,t')\wedge\phi[t']\}} p(t, s')}{\sum_t p(t, s')}$$

and hence, by definition of $BelThatInPrev$, it holds that $BelThatInPrev(\phi, s') = x$, as we wanted to show. ∎

3.6 Example

Let us show how the results of the previous section formalize, in a concrete scenario, the following concepts:

- An agent starting off with a reasonably high confidence in a formula, and being able to introspect said confidence with full certainty.

- Said agent carrying out a sensing action which causes its degree of belief in that same formula to drop significantly.

- The agent having full certainty of its new degree of belief in the formula as well as its previous one, and consequently, full certainty in its degree of belief change in said formula.

Let us assume we have an agent that has beliefs regarding its distance in meters to the nearest wall. The value of this distance will be represented by the fluent *distance*.

Now, let us consider a formula ϕ which will hold if and only if the agent's distance to the nearest wall in the current situation is between 50 and 150 meters. In other words:

$$\phi \stackrel{\text{def}}{=} distance(now) = 50 \vee distance(now) = 51 \vee \cdots \vee distance(now) = 150$$

Let us assume that, initially, the agent's beliefs about distance are distributed according to a Gaussian of mean 100 and deviation 30, plus a small uniform, i.e. we assume the following holds:

$$\forall x. Bel(distance(now) = x, S_0) \propto Normal((x - 100)/30) + 0.0001$$

For this example we will assume that the fluent *distance* can only take integer values between 0 and 1000.

The agent's distribution of beliefs for *distance*, according to $Bel(distance(now) = x, S_0)$ can be seen on the graphic at the top of figure 3.1 plotted as a solid line. The agent's degree of belief in formula ϕ corresponds to the shaded area under the curve. According to the agents belief distribution at S_0, it will hold that $Bel(\phi, S_0) = 0.835$. Also note that, as a consequence of theorem 3.2.1, it will also hold that $Bel(Bel(\phi, now) = 0.835, S_0) = 1$.

Now, let us assume that the agent senses its distance to the nearest wall by means of the action $\mathsf{noisy\text{-}sense\text{-}distance}$ defined as:

$$\mathsf{noisy\text{-}sense\text{-}distance} \stackrel{\text{def}}{=} \pi x, y.\mathsf{sense\text{-}distance}(x, y)$$

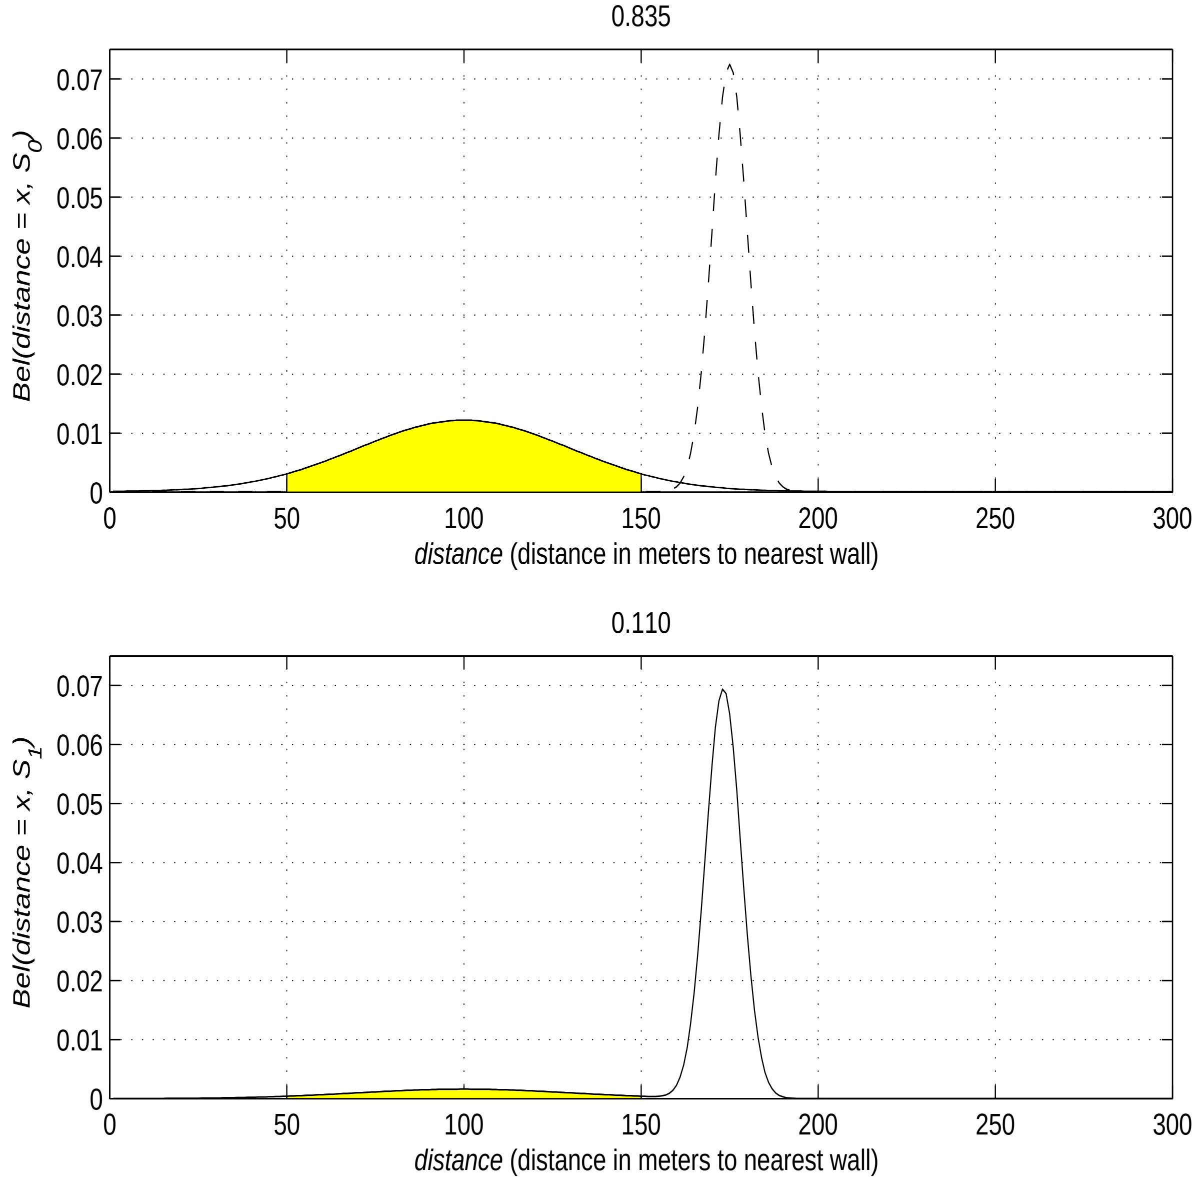

Figure 3.1: The x axis represents the agent's distance in meters to the wall. **Top**: The solid curve plots $Bel(distance(now) = x, S_0)$. The dashed curve shows the likelihood values that will be used to calculate p for situation S_1 **Bottom**: The solid curve plots $Bel(distance(now) = x, S_1)$. In both cases the shaded area under the solid curve represents the agent's degree of belief in formula ϕ: "the distance to the wall is between 50 and 150 meters."

We will define the likelihood function ℓ for $\mathsf{sense\text{-}distance}$ as a Gaussian function of x of mean y and deviation 10 plus a small uniform. Formally:

$$\ell(\mathsf{sense\text{-}distance}(x, y), s) \propto \begin{cases} Normal((x - y)/10) + 0.0001 & \text{if } y = distance(s) \\ 0 & \text{if } y \neq distance(s) \end{cases}$$

Let's see how the agent's beliefs change when it senses a distance value of 200 meters. Let $S_1 = do(\mathsf{sense\text{-}distance}(200, f(S_0)), S_0)$.

The top of figure 3.1 shows the likelihood values that will be used to calculate the new belief distribution plotted as a dashed line centered around the sensed value 200.

If we now compute the agent's distribution of beliefs for *distance* in situation S_1, we obtain the belief curve shown at the bottom of figure 3.1. As before, the shaded area represents the degree of belief in formula ϕ. According to the updated values, it will now hold that $Bel(\phi, S_1) = 0.110$, and, as before, as a consequence of theorem 3.2.1 it will also hold that $Bel(Bel(\phi, now) = 0.110, S_1) = 1$.

As for past beliefs, in S_1 it now holds that $PrevBel(\phi, S_1) = 0.835$; *i.e.* in situation S_1, the agent's degree of belief in formula ϕ in the previous situation was 0.835. From theorem 3.3.1 we also can conclude that $Bel(PrevBel(\phi, S_1) = 0.835) = 1$ holds; in other words, the agent has full certainty that in the situation previous to S_1 it believed ϕ with a degree of 0.835 and now, in situation S_1, it believes ϕ with a degree of 0.110.

From these results, we can conclude that $BelChange(\phi, S_1) = -0.725$ holds, and through corollary 3.4.1 we also know that it holds that $Bel(BelChange(\phi, S_1) = -0.725) = 1$. In other words, in S_1, the agent has full certainty that from the previous situation to the present one, its degree of belief in ϕ dropped by 0.725.

Now, let us show how, in this scenario, the agent can come to believe that its previous beliefs about ϕ were mistaken.

Given that we know that the performed sensing action does not change the agent's distance to the wall, the value of $BelThatInPrev(\neg\phi, S_1)$ is the same as $Bel(\neg\phi, S_0)$, which in turn evaluates to the same value as $1 - Bel(\phi, S_0)$. We may then conclude that it holds that $BelThatInPrev(\neg\phi, S_1) = 0.8904$. So the agent now believes with degree 0.8904 that in the previous situation ϕ was false.

Also, we know that it holds that $Bel(\phi, S_0) = 0.835$, and using the same reasoning we used in our proof of theorem 3.2.1 we may conclude that $Bel(\phi, s') = 0.835$ holds in all situations p-related to S_0, which are in turn the situations preceding situations p-related to S_1. This obviously includes all situations where $\neg\phi$ holds. Then, by definition of $BelThatInPrev$ we may conclude that it holds that

$$BelThatInPrev(Bel(\phi, now) = 0.835 \wedge \neg\phi, S_1) = 0.8904$$

This is interpreted as the agent believing with degree 0.8904 that it believed with degree 0.835 in a false formula. Finally, by theorem 3.5.1 the agent will have full introspection of this fact. In other words, it holds that

$$Bel(BelThatInPrev(Bel(\phi, now) = 0.835 \wedge \neg\phi, S_1) = 0.8904) = 1.$$

Chapter 4

Sensor coarseness

4.1 Coarseness vs. noise

In [1], it is shown how the Bacchus *et al.* extension to the situation calculus may be used to represent different forms of noisy sensors. This is taken to mean, sensors which are expected to exhibit certain inaccuracies in their readings.

In this section however, we deal with another aspect of sensor reading inaccuracy, an aspect we will refer to as the *coarseness* of the sensor. In other words, the actual measurement units the sensor is using.

For example, let us assume we have an agent equipped with a sensor that provides it with information as to its distance to a wall. In order of increasing coarseness, the sensor's readings might be millimeters, centimeters, meters, decameters, *etc.* When equipped with a sensor of millimetric coarseness, the agent would receive values such as 112.578 meters. In other words, it would receive information up to the millimeter. When equipped with a sensor capable only of obtaining readings to the nearest meter, the agent, being at the same distance from the wall, would receive a value rounded to the nearest meter, in this case 113 meters. Alternatively a truncated value of 112 meters is also possible, depending on the characteristics of the sensor.

4.2 Modeling varying coarseness

Given a fixed level of coarseness for two sensors, they may have different noise characteristics: *i.e.* one may be more reliable than the other. This is clearly modeled in the framework by letting us describe different likelihood distributions between a sensor's obtained value, and the true value of the fluent being sensed. Less apparent is the fact that the dual of this is also true. In other words, two sensors may have basically the same likelihood function over values read and true values while having different coarseness.

We will now show how, assuming we have likelihood distribution regarding the sensor's raw signal, i.e. previous to any rounding due to coarseness, we may obtain the likelihood distribution for the coarse values.

Let us call the fluent being sensed f, the value sensed x and the fluent's true value $y = f(s)$ for some situation s. The action-likelihood function for the sensing action sense for the fluent f will give us the aforementioned likelihood distribution regarding the sensor's raw signal. This action, however, will not be available to the agent. We use it here only to obtain the action-likelihood function for the coarse sensing.

We will also need a coarseness value c for the sensor. This corresponds to the unit distance between consecutive sensor values. For example if the sensor reading corresponds to meters from the agent to the wall and the values returned are rounded to the nearest 100 meters, then we would have $c = 100$.

Given all these elements, we will associate with our coarse sensor a sensing action $\mathsf{sense\text{-}coarse}_c$, with its action-likelihood function defined as follows

$$\ell(\mathsf{sense\text{-}coarse}_c(x, y), s) \propto \int_{x-c/2}^{x+c/2} \ell(\mathsf{sense}(x', y), s)\, dx' \tag{4.1}$$

This gives us the likelihood of y being the true value of f given that the value x was sensed using the coarse sensor.

Perhaps this expression can be better understood by means of an example: assume, once more, that f corresponds to distance in meters to a wall, $c = 100$, $x = 100$, and $y = 112$. Since any more accurate (to the meter) reading between 50 and 150 would be rounded to a coarse reading of 100, the likelihood of 112 being the true value given that the value sensed is 100 is proportional to the sum of the likelihoods of 112 being the true value given that we sensed anything between 50 and 150 in the uncoarse model.

Let us also take the opportunity to make note of the fact that $\textsf{sense-coarse}_c$ should only be possible for values of x (the values corresponding to those returned by the sensor) are divisible by c. For example, still assuming $c = 100$, $\textsf{sense-coarse}_c(100, 112)$, $\textsf{sense-coarse}_c(200, 112)$, and $\textsf{sense-coarse}_c(300, 112)$ are all possible. However actions like $\textsf{sense-coarse}_c(101, 112)$ should not be considered, for one cannot receive a sensor value of 101 with a coarseness of $c = 100$.

We shall give an example of how this previous formula can be represented logically. Suppose that the correlation between the sensed value and the true value, x and y respectively, is described by a linear Gaussian model. Then, we may instantiate formula 4.1 into logical form as follows:

$$\ell_{un-norm}(\textsf{sense-coarse}_c(x, y), s) = Normcdf(\frac{x + c/2 - y}{\sigma}) - Normcdf(\frac{x - c/2 - y}{\sigma}) \tag{4.2}$$

$$\ell(\textsf{sense-coarse}_c(x, y), s) = \frac{\ell_{un-norm}(\textsf{sense-coarse}_c(x, y), s)}{\sum_{y'} \ell_{un-norm}(\textsf{sense-coarse}_c(x, y'), s)} \tag{4.3}$$

The right hand side of formula 4.2 corresponds to the right hand side of formula 4.1. In order to compute the integral required, we take the difference of said integral evaluated at the two integration extremes. $Normcdf$ corresponds to the cumulative distribution function of a standardized Gaussian. Similarly to what we've done before with $Normal$, in order to allow us to compute it logically, we assume that given some fixed precision, a table of values for $Normcdf$ is computed and added to the domain theory as a set of equations.

Formula 4.3 normalizes the values of $\ell_{un-norm}$. In order to make the sum in the denominator finite, the agent must have a finite and discrete set of values that it will consider as possible true values for the fluent being sensed. For instance, if the fluent being sensed is distance in meters, and the agent distinguishes among beliefs differing in at least 1 centimeter (0.01 meters), and assuming it considers a maximum distance between objects of 1 kilometer (1000 meters), the set would consist of $\{0, 0.01, 0.02, ..., 1000\}$. In an appendix of [1], Bacchus *et al.* show how a summations can be formalized in the situation calculus.

Figure 4.1 illustrates likelihood curves of increasing coarseness. These are generated using equation 4.1. The action-likelihood function for action sense is a Gaussian of variance 10. The first (tallest) curve is the one corresponding to a sensor coarseness equal to the coarseness of the beliefs handled by the agent: 1 meter intervals. In the following curves, the coarseness increases by powers of 10: the second tallest curve corresponds to a coarseness of 10 meter intervals, the third to 100 meter intervals, and finally, the last one corresponds to 1000 meter intervals. What becomes apparent is that, as the coarseness increases, the curves become more uniform around 0. This corresponds to situations where the value sensed and the true value differ in less than half the sensor coarseness. Given that when we perform Bayesian updating using a uniform, the updated curve remains unchanged, this suits our purposes modeling the fact that little information should be gained from such a sensing.

4.3 Sensor coarseness and belief update

Intuitively, the greater the coarseness of the sensor being used, the less information an agent will be receiving from sensed values. For example, using the same scenario as before, let's say an agent currently believes with highest probability that it is 112 meters away from the nearest wall. One would expect the agent's beliefs not to change significantly

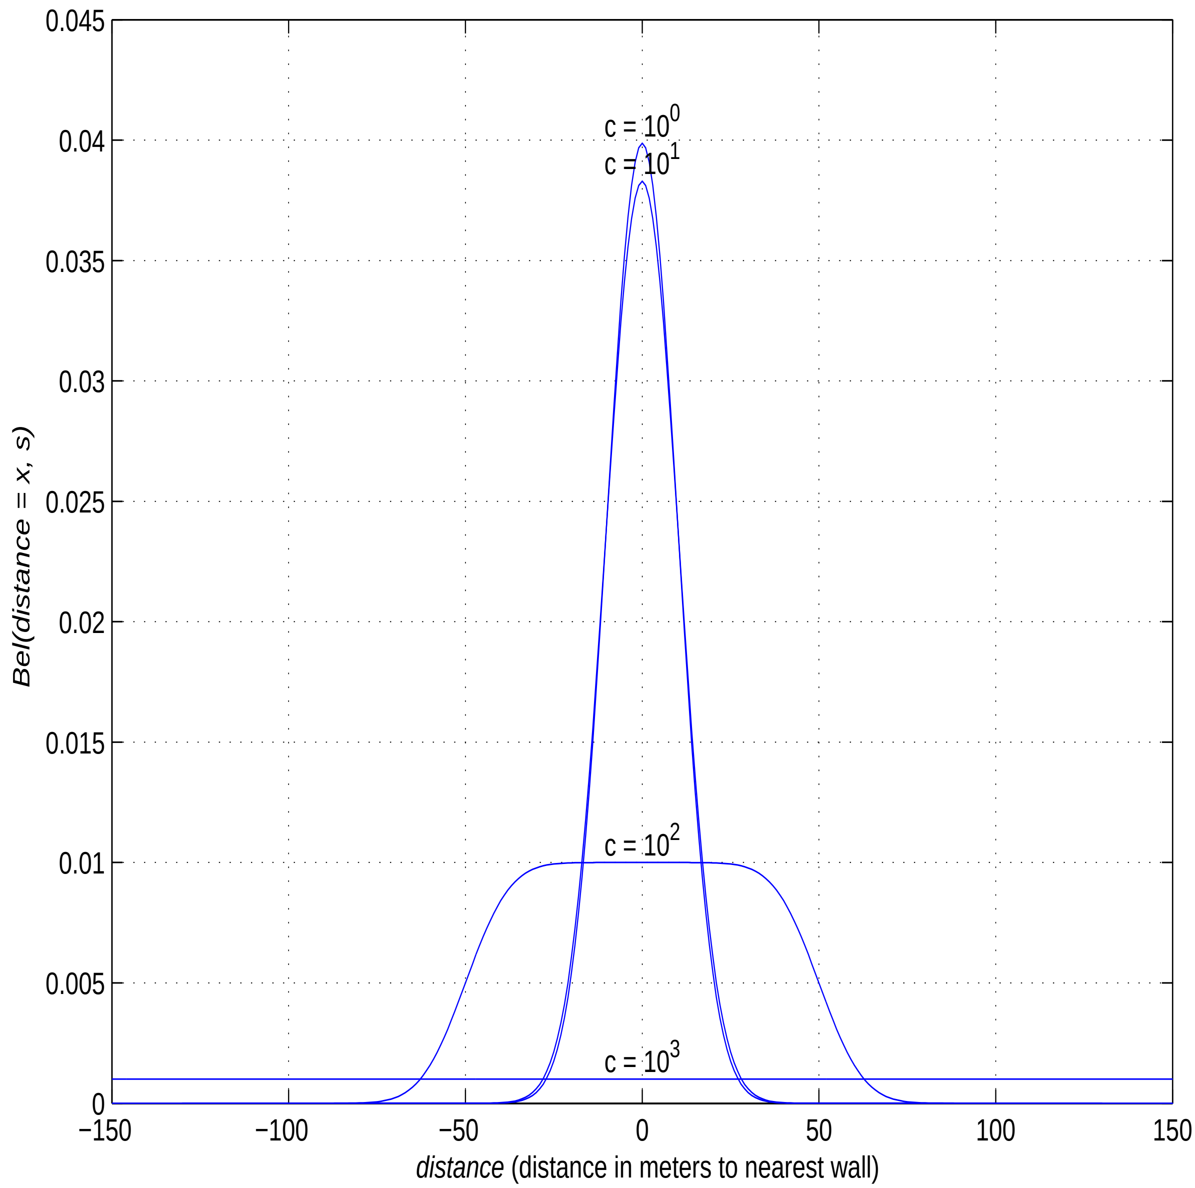

Figure 4.1: Likelihood distributions for sensors of varying coarseness using a fixed normal distribution to model noise.

after receiving repeated readings of 100 meters, if the sensor being used has a coarseness of 100 meter intervals.

The belief update model takes into account that the granularity of the sensing might be coarser than the agent's beliefs [1]. The coarser the sensor readings (i.e. the less "significant digits" the sensor has), the less the agent's beliefs should change as a result of the sensing. This is reflected by the fact the agent is learning less through the sensing.

4.4 Example

Let us illustrate the matter by an example.

Let us assume the agent's prior beliefs are centred around 112 meters, and normally distributed with a standard deviation of 10 meters. In other words, it holds that

$$\forall x.Bel(distance(now) = x, s) = Normal((x - 112)/10)$$

This can be seen in graphic form in figure 4.2, in the top left hand corner. Let us further assume that the true distance to the wall is of 100 meters.

Now, let us first equip the agent with a sensor of coarseness of 1 meter intervals whose noise is modeled by a Gaussian of a standard deviation of 10 meters.

Figure 4.2 shows how the agent's beliefs change after repeated sensings of values distributed according to the previously mentioned likelihood. As can be seen, after five sensings the belief curve has sharpened around the true value of 100 meters.

Now, let us see what happens if we change the coarseness of the sensor to that of 100 meter intervals, while keeping the underlying noise modeling function fixed. Once again, we begin with the agent's beliefs centred around 112 meters, and normally distributed with a standard deviation of 10 meters. Figure 4.3 shows this graphically in the top left

[1] Note that the reverse of his situation would be to have a sensor who's return values are finer grained than the agent's beliefs allow for, in which case, they'd have to be rounded after update has taken place.

hand corner.

Figure 4.3 also shows how the updates take place now that the coarseness has been increased. As can be seen, even after the five sensings of values distributed accordingly with the sensor's noise and coarseness, the agent's initial beliefs remain virtually unchanged. This is because most of the agent's (noticeable) beliefs are allocated to values that fall under the uniform segment of the likelihood curve used to update said beliefs. If the agent would have had non-zero values allocated to distances greater than 150, for example, or lower than 50, these would have been lowered considerably after each successive sensing result. This corresponds to the reasoning that even though the sensor is coarse, the agent is more convinced that the true value is in the range $(x - c/2, x + c/2)$, in this case, $(50, 150)$.

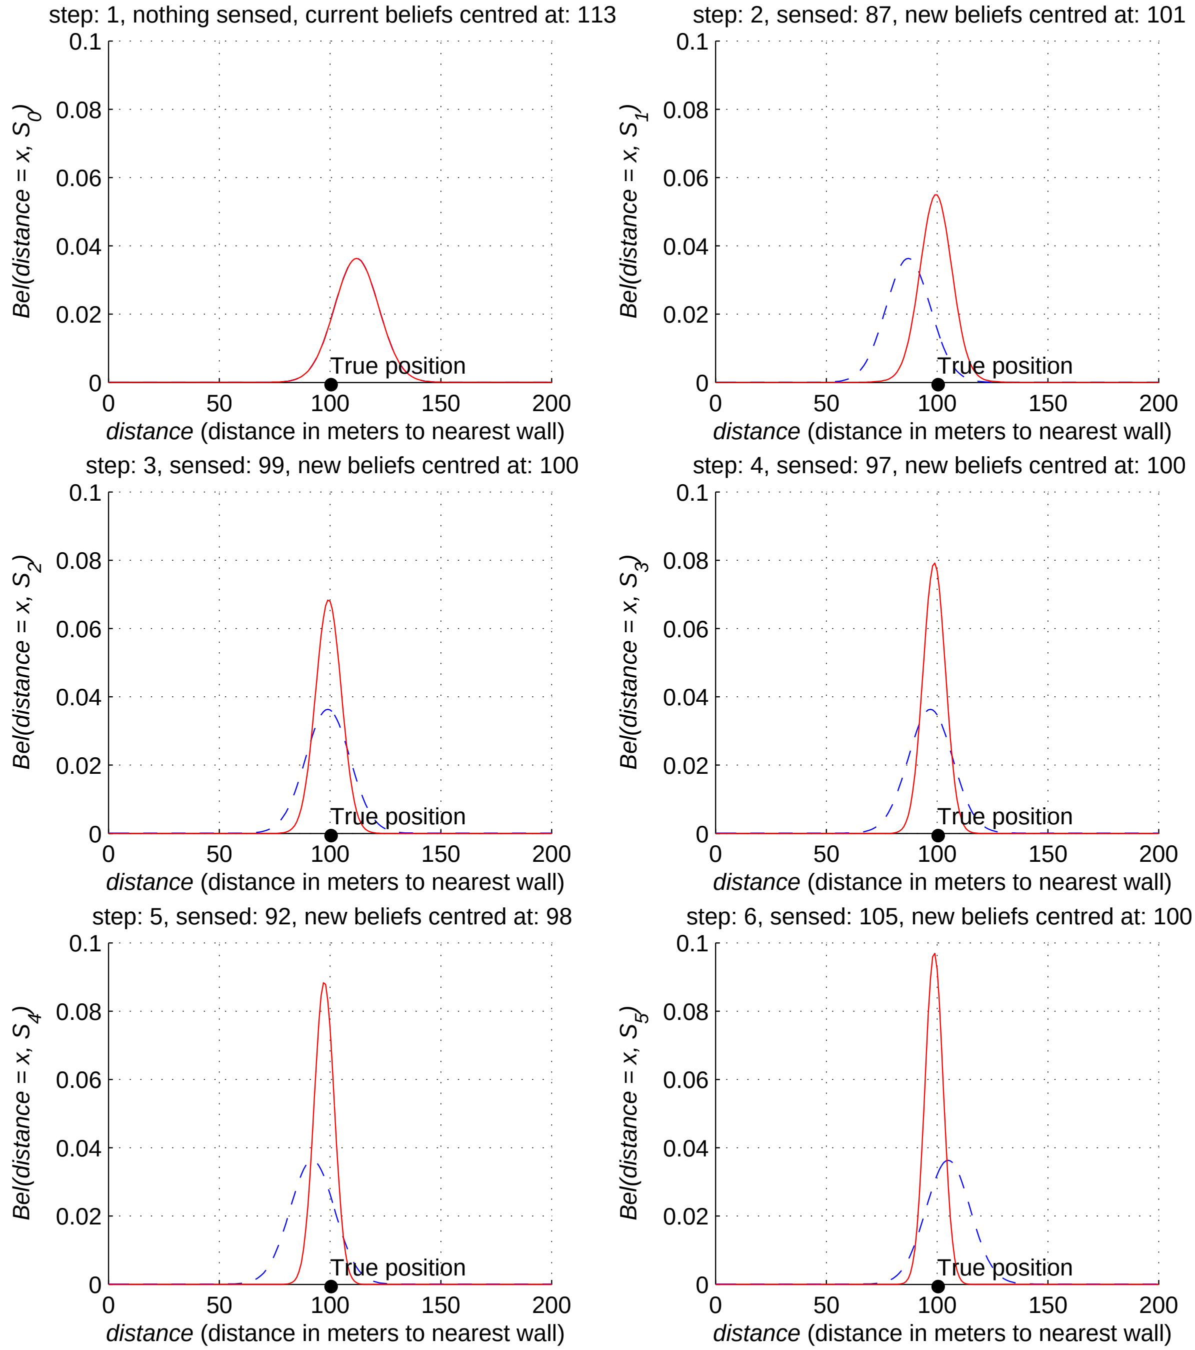

Figure 4.2: Solid curve represents the agent's current belief distribution as to its distance to the wall. Dashed curve represents the likelihood distribution used in this step to obtain said current beliefs from the previous ones. The agent is equipped with a sensor of coarseness of 1 meter intervals whose noise is modeled by a Gaussian of a standard deviation of 10 meters. The true distance to the wall is 100 meters.

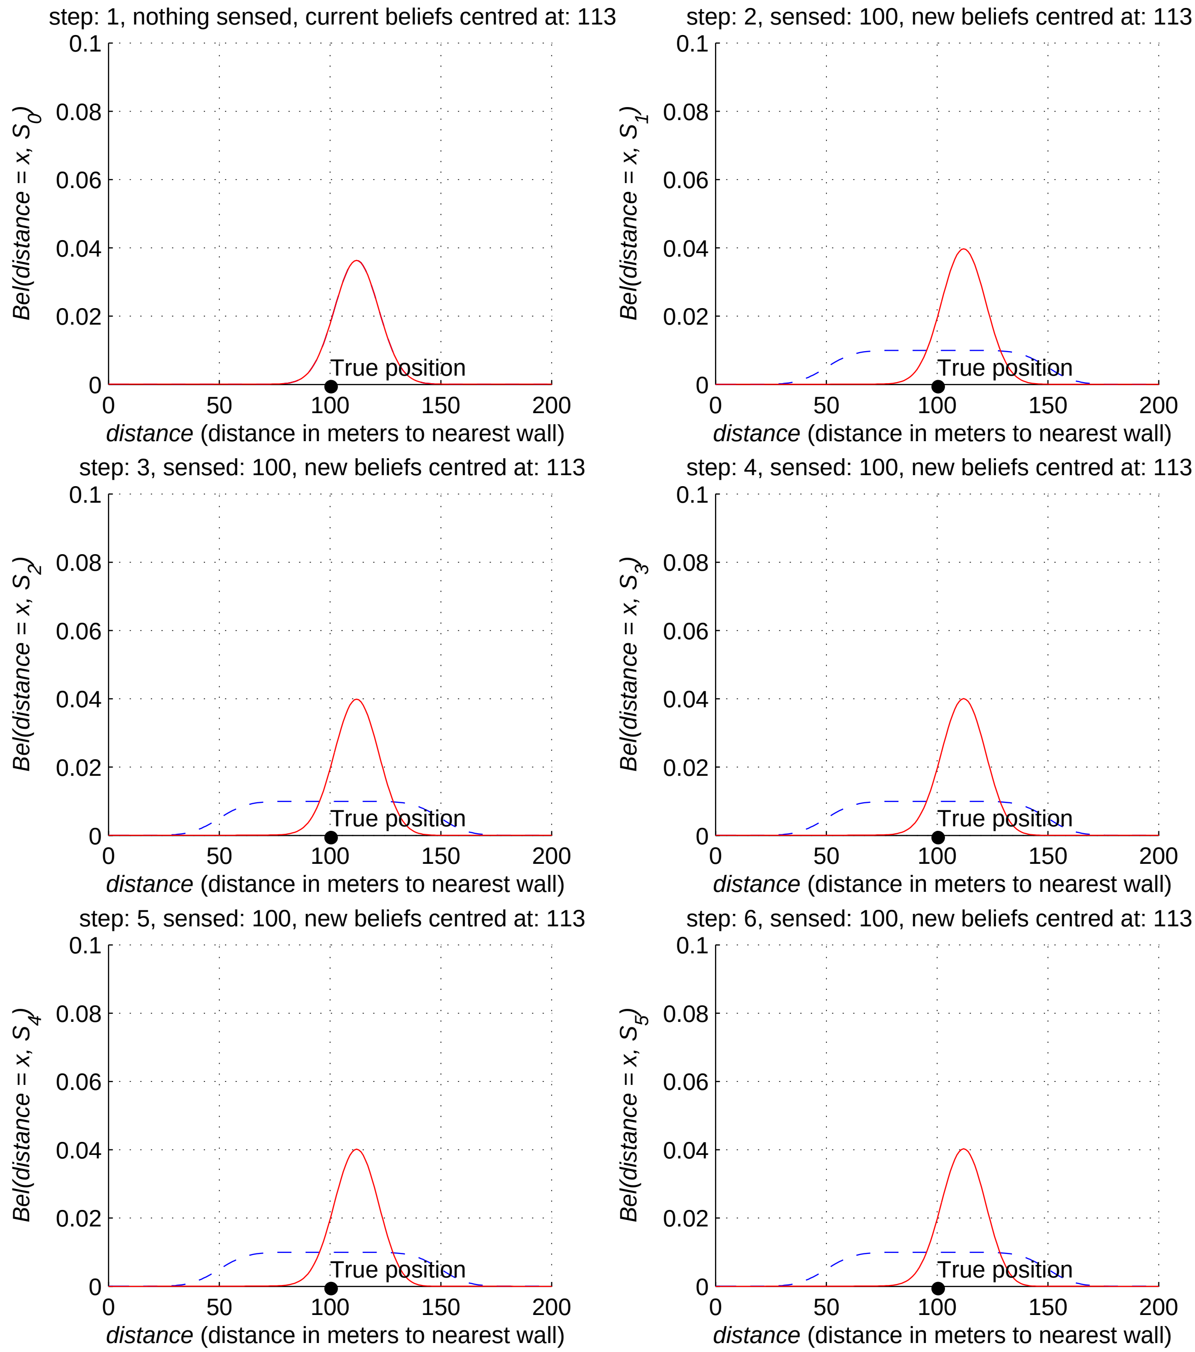

Figure 4.3: Solid curve represents the agent's current belief distribution as to its distance to the wall. Dashed curve represents the likelihood distribution used in this step to obtain said current beliefs from the previous ones. Here, the agent is equipped with a sensor of coarseness of 100 meter intervals whose noise is modeled by a Gaussian of a standard deviation of 10 meters. The true distance to the wall is 100 meters.

Chapter 5

Combining noisy sensors: Conditional probability densities and the Bacchus *et al.* framework

In this chapter, we will show by example how the Bacchus *et al.* framework can be used in such a way as to be equivalent to using conditional probability densities for combining noisy sensors.

5.1 Combining readings using conditional probability densities

In many real world situations, an agent may have access to multiple sensors, all of which provide it with a reading regarding the same variable whose value the agent wishes to estimate. In [13], we find the following intuitive example of such a scenario: an agent wishes to estimate the velocity of an aircraft. For such a purpose, it has access to

- a Doppler radar

- the velocity indications of an inertial navigation system
- the pitot and static pressure and relative wind information in the air data system.

all of which provide the agent with an estimate of the aircraft's velocity.

In a scenario such as this one, it is desirable to use all available sensor information along with what we know about the system's dynamics and combine it into an estimate that is better than any of the individual sensor readings considered separately. From a Bayesian perspective, we want to "propagate the conditional probability density of the desired quantities, conditioned on knowledge of the actual data coming from the measuring devices" ([13]). In essence, the agent's beliefs after sensing will be represented as a distribution whose mode will be the agent's estimate for the variable's value. Furthermore, we wish this estimate to have maximum a posteriori value given our observations and what is known about sensor noise.

As an example, let us show how we may use conditional probability densities to combine two sensor readings subject to zero-mean white Gaussian noise.

Assume the agent wishes to estimate the value of a variable x, for which it has uniform prior. Furthermore, assume that the first and second sensor have associated a standard deviation of σ_1 and σ_2 respectively, and that they provide the agent with independent readings of $x_1 = \mu_1$ and $x_2 = \mu_2$. Thus, we have

$$p(x|x_1 = \mu_1) = \mathcal{N}(x; \mu_1, \sigma_1)$$

$$p(x|x_2 = \mu_2) = \mathcal{N}(x; \mu_2, \sigma_2)$$

We calculate the conditional probability distribution based on both readings as follows:

$$p(x|x_1 = \mu_1, x_2 = \mu_2) \propto p(x_1 = \mu_1, x_2 = \mu_2|x)p(x)$$

Given that we are assuming that the agent has a uniform prior regarding x, we may obtain

$$p(x|x_1 = \mu_1, x_2 = \mu_2) \propto p(x_1 = \mu_1, x_2 = \mu_2|x)$$

But since the readings are conditionally independent given x, we have

$$p(x|x_1 = \mu_1, x_2 = \mu_2) \propto p(x_1 = \mu_1|x)p(x_2 = \mu_2|x)$$

Furthermore,

$$\begin{aligned} p(x_1 = \mu_1|x) &\propto p(x|x_1 = \mu_1)p(x_1 = \mu_1) \\ &\propto p(x|x_1 = \mu_1) \\ &= \mathcal{N}(x; \mu_1, \sigma_1) \end{aligned}$$

Analogously, we obtain

$$p(x_2 = \mu_2|x) \propto \mathcal{N}(x; \mu_2, \sigma_2)$$

Replacing above, we now have

$$p(x|x_1 = \mu_1, x_2 = \mu_2) \propto \mathcal{N}(x; \mu_1, \sigma_1)\mathcal{N}(x; \mu_2, \sigma_2)$$

Since we know that the product of two Gaussians $\mathcal{N}(\mu_1, \sigma_1)$ and $\mathcal{N}(\mu_2, \sigma_2)$ is proportional to a third Gaussian, whose mean μ and standard deviation σ are given by

$$\begin{aligned} \mu &= [\sigma_2^2/(\sigma_1^2 + \sigma_2^2)]\mu_1 + [\sigma_1^2/(\sigma_1^2 + \sigma_2^2)]\mu_2 \\ 1/\sigma^2 &= (1/\sigma_1^2) + (1/\sigma_2^2) \end{aligned}$$

we may arrive at

$$p(x|x_1 = \mu_1, x_2 = \mu_2) \propto \mathcal{N}(x; \mu, \sigma)$$

Hence, the maximum a posteriori estimate for the value of the variable in question given said sensor readings is μ, the mean of the Gaussian distribution proportional to the product of the Gaussian distributions $\mathcal{N}(\mu_1, \sigma_1)$ and $\mathcal{N}(\mu_2, \sigma_2)$.

The following section shows an example of how we would use this formulation in a scenario in which we wish to combine the readings obtained from two noisy sensors.

5.2 Example

Let's assume that an agent wishes to have a good approximation of its distance the nearest wall. For the purposes of determining such a distance, it is equipped with two sensors, both of which return a value in meters. These sensors are subject to zero-mean white Gaussian noise. The first sensor has a standard deviation of $\sigma_1 = 100$ meters, while the second has a standard deviation of $\sigma_2 = 50$ meters. We will assume that the agent's prior is uniform over all possible distances.

Upon using the first sensor, the agent reads a value of 400 meters. At this point, the conditional probability of the agent's distance to the wall, conditioned on this first value read, will be a Gaussian curve of mean $x_1 = 400$ meters (the value read) and standard deviation of $\sigma_1 = 100$ meters (the uncertainty associated with this particular sensor.) This curve is shown on the top left of figure 5.1.

Now the agent uses its second sensor and receives a reading of 800 meters. Using the information obtained from this second reading alone, the conditional probability distribution of the agent's distance to the wall will be a Gaussian of mean $x_2 = 800$ meters and a standard deviation of $\sigma_2 = 50$ meters. This second distribution is plotted on the top right of figure 5.1. For comparison purposes, the distribution associated with the first reading is plotted as a dashed line.

Now, the agent may combine both readings. The conditional probability distribution of the agent's position given *both* readings x_1 and x_2 will be a Gaussian of mean μ and standard deviation σ where:

$$\begin{aligned}\mu &= [\sigma_2^2/(\sigma_1^2+\sigma_2^2)]x_1 + [\sigma_1^2/(\sigma_1^2+\sigma_2^2)]x_2 = \\ &= [50^2/(100^2+50^2)]400 + [100^2/(100^2+50^2)]800 = \\ &= 720 \text{ meters}\end{aligned}$$

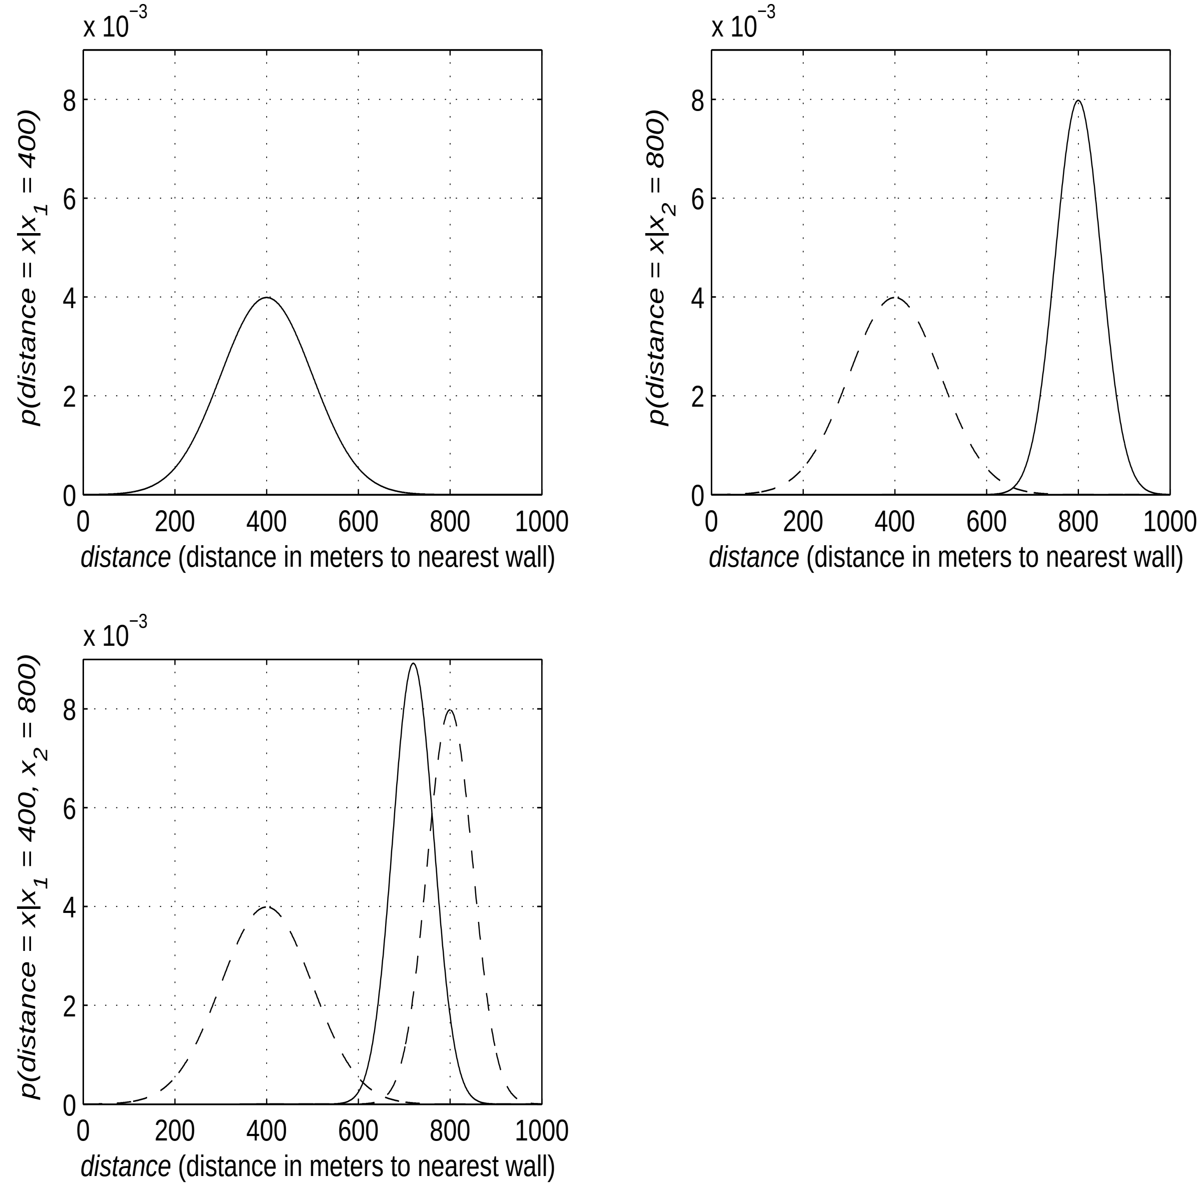

Figure 5.1: **Top left**: Conditional probability distribution based on the agent's first reading of 400 meters. **Top right**: Conditional probability distribution based on the agent's second reading of 800 meters alone. For comparison purposes, the distribution associated with the first reading is plotted as a dashed line. **Bottom left**: The conditional probability distribution of the agent's position given both readings. For comparison purposes the distributions associated with the first and second readings alone are plotted as a dashed line.

$$\begin{aligned}
\sigma &= [(1/\sigma_1^2) + (1/\sigma_2^2)]^{-\frac{1}{2}} \\
&= [(1/100^2) + (1/50^2)]^{-\frac{1}{2}} \\
&\approx 44.72136 \text{ meters}
\end{aligned}$$

This new distribution is plotted at the bottom left of figure 5.1. As before, for comparison purposes, the distributions associated with the first reading and second readings alone are plotted as a dashed line.

Given both sensor readings, the agent's best estimate of its distance to the nearest wall will be $\mu = 720$ meters. It should also be noted that the this probability distribution is narrower around its mean, indicating that the agent's uncertainty about its distance is less than with either of the two readings considered independently.

We should point out here that the reason the agent's best estimate converges to 720 meters after two wildly different readings of 400 and 800 meters is because of the very strong assumption that the sensors are truly Gaussian. In real-world scenarios, most sensors will be modeled as a mixture of a Gaussian and a much more heavily tailed distribution, or even a uniform. (We showed examples of this in chapters 2 and 3.)

5.3 Comparison with the Bacchus *et al.* framework

Let us show how the Bacchus *et al.* framework can be equivalent to using conditional probability densities to obtain a maximum a posteriori probability estimate as in the previous section.

Suppose we are dealing with an agent equipped with two sensors which read the value of some fluent f. Furthermore, let us assume that these sensors are modeled as having zero-mean white Gaussian noise, the first one with standard deviation σ_1 and the second with standard deviation σ_2, as we did in the example of the previous section.

In the Bacchus*et al.* framework, we model this scenario by defining two sensing actions associated with each of the sensors. Let us call these actions $\mathsf{noisy\text{-}sense\text{-}f_1}$ and $\mathsf{noisy\text{-}sense\text{-}f_2}$ defined as:

$$\mathsf{noisy\text{-}sense\text{-}f_1} \stackrel{\text{def}}{=} \pi x, y.\mathsf{sense\text{-}f_1}(x, y)$$
$$\mathsf{noisy\text{-}sense\text{-}f_2} \stackrel{\text{def}}{=} \pi x, y.\mathsf{sense\text{-}f_2}(x, y)$$

Given what we know about the sensors, we define the action-likelihood functions for each of these actions as follows:

$$\ell(\mathsf{sense\text{-}f_1}(x, y), s) = \begin{cases} Normal((y - x)/\sigma_1) & \text{if } y = f(s) \\ 0 & \text{if } y \neq f(s) \end{cases}$$

$$\ell(\mathsf{sense\text{-}f_2}(x, y), s) = \begin{cases} Normal((y - x)/\sigma_2) & \text{if } y = f(s) \\ 0 & \text{if } y \neq f(s) \end{cases}$$

Now, assume the agent uses the first sensor and reads a value of x_1 giving us $S_1 = do(\mathsf{sense\text{-}f_1}(x_1, f(S_0)), S_0)$.

According to the successor state axiom for p [1], for all situations s'^{+} and s' such that $s'^{+} = do(\mathsf{sense\text{-}f_1}(x_1, f(s')), s')$ it will hold that

$$p(s'^{+}, S_1) = p(s', S_0) \times \ell(\mathsf{sense\text{-}f_1}(x_1, f(s')), s')$$

Now assume that, subsequently, the agent uses the second sensor, reading a value of x_2, giving us $S_2 = do(\mathsf{sense\text{-}f_2}(x_2, f(S_1)), S_1)$. Now we will have that for all situations s'^{+} and s' such that $s'^{+} = do(\mathsf{sense\text{-}f_2}(x_2, f(s')), s')$ it will hold that

[1] Notice that we are now using p as defined in chapter 2 in the context of the Bacchus *et al.* framework.

$$p(s'^{+}, S_2) = p(s', S_1) \times \ell(\textsf{sense-f}_2(x_2, f(s')), s')$$

By combining the previous two equalities we may conclude that for all situations s'^{+} and s' such that $s'^{+} = do(\textsf{sense-f}_2(x_2, f(s')), do(\textsf{sense-f}_1(x_1, f(s')), s'))$ it will hold that

$$p(s'^{+}, S_2) = p(s', S_0) \times \ell(\textsf{sense-f}_1(x_1, f(s')), s') \times \ell(\textsf{sense-f}_2(x_2, f(s')), s')$$

Notice here that we are multiplying p by the product of the two action-likelihood functions which are Gaussians of mean x_1 and x_2 and standard deviation σ_1 and σ_2 respectively. The result is proportional to multiplying directly by one likelihood function corresponding to a Gaussian of mean x_3 and standard deviation σ_3 where

$$x_3 = [\sigma_2^2/(\sigma_1^2 + \sigma_2^2)]x_1 + [\sigma_1^2/(\sigma_1^2 + \sigma_2^2)]x_2$$

$$1/\sigma_3^2 = (1/\sigma_1^2) + (1/\sigma_2^2)$$

5.4 Example continued

In order to illustrate the equivalency we have just shown above, we will show the behaviour of the Bacchus *et al.* framework under the same scenario as we used in the conditional probability densities example.

The fluent f being sensed will be an integer value representing the agent's distance, in meters, to the nearest wall. We will assume that there are two sensing actions associated with this fluent corresponding to each of the sensors. These actions and their associated action-likelihood functions are defined as discussed in the previous subsection.

For simplicity, we will assume the agent's beliefs in the initial situation S_0 regarding f are uniform, as shown on the top left of figure 5.2 (as we did in the conditional probability densities example.)

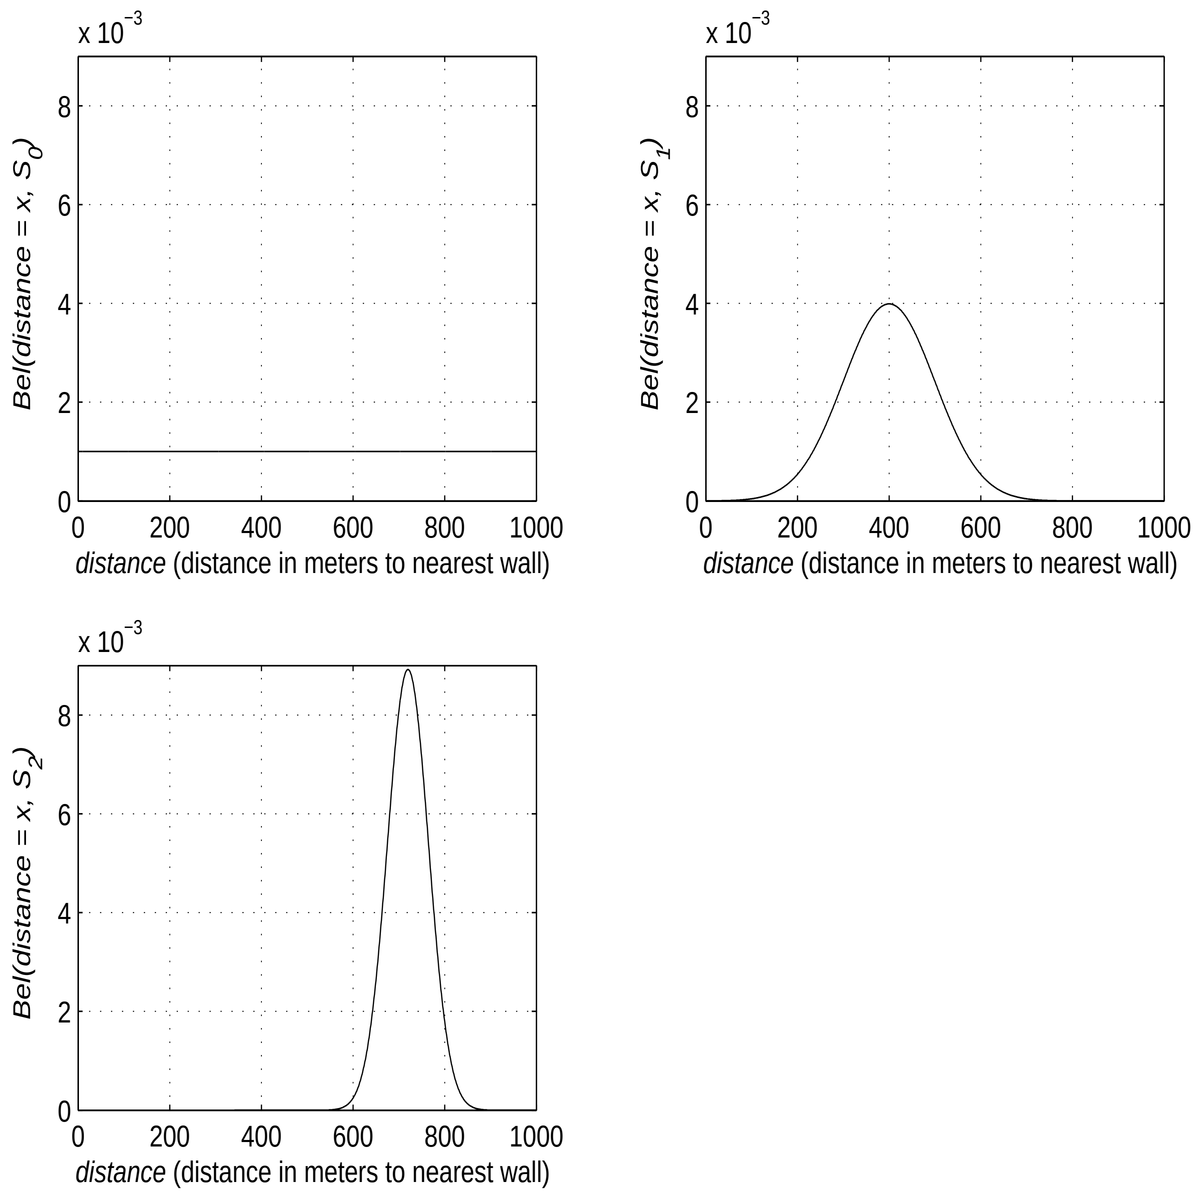

Figure 5.2: **Top left**: In situation S_0, the agent's beliefs regarding f are uniformly distributed. **Top right**: Distribution for $Bel(f(now) = x, S_1)$, where $S_1 = do(\mathsf{sense\text{-}f_1}(400, f(S_0)), S_0)$ for x varying from 0 to 1000. **Bottom right**: Distribution for $Bel(f(now) = x, S_2)$, where $S_2 = do(\mathsf{sense\text{-}f_2}(800, f(S_0)), S_1)$, x varying from 0 to 1000.

First, we model the fact that the agent senses its distance to the nearest wall using the first sensor and obtains a value of 400. This is done by saying that the agent performed the action $\mathsf{noisy\text{-}sense\text{-}f}_1$. The graph on the top left of figure 5.2 plots $Bel(f(now) = x, S_1)$, where $S_1 = do(\mathsf{sense\text{-}f}_1(400, f(S_0)), S_0)$ for x varying from 0 to 1000.

Next, the agent uses its second sensor and obtains a value of 800. This is modeled by the agent performing the action $\mathsf{noisy\text{-}sense\text{-}f}_2$. The resulting belief distribution for $Bel(f(now) = x, S_2)$, where $S_2 = do(\mathsf{sense\text{-}f}_2(800, f(S_1)), S_1)$, x varying from 0 to 1000, is plotted on the bottom left of figure 5.2.

As expected from what was shown in the previous section, the resulting curve is identical to the solid curve of at the bottom left figure 5.1 corresponding to the conditional probability distribution of the agent's position given both readings.

We should note that in the case of the Bacchus *et al.* framework we are dealing with discrete beliefs, were as when we were dealing with conditional probabilities, the distribution was continuous. For the results to be truly identical, we should first discretize the continuous conditional probability distribution we obtained to match the granularity of beliefs we've chosen for the agent modeled with the Bacchus *et al.* framework.

Chapter 6

Conclusions

We have seen how the framework of Bacchus *et al.*, an extension of the situation calculus created to model noisy actions, while dealing with probabilistic beliefs, can also model concepts such as belief introspection, previously held beliefs, current beliefs about the past, awareness of belief change and belief in the fact that previously held beliefs were mistaken. Moreover, we proved that in this framework, under certain conditions, the agent will have full introspection of its beliefs. Similar results were shown for the introspection of previously held beliefs, current beliefs about the past and belief change.

We also proved that, when a sensor is modeled as being noise free, the framework is well behaved with regards to belief, resulting in the agent having full certainty of the sensed value.

In addition, we addressed the problem of sensor coarseness, showing how the framework of Bacchus *et al.* can be adapted to model scenarios in which it is possible for an agent to be equipped with a sensor that returns values which are coarser than the agent's beliefs. We also discussed the consequences of coarseness for belief update.

Finally, we compared the framework of Bacchus *et al.* to combining noisy sensor data using conditional probability densities. We showed that the Bacchus *et al.* framework can function in a way which is equivalent.

As for future work, given characteristics of the likelihood function associated with the sensing action of a given fluent, it would be interesting to prove bounds on the number of repeated sensings necessary before an agent becomes “reasonably convinced” of a fluent’s true value. In other words, given some fluent f and small real number ϵ, how many sensings does the agent have to carry out before it holds that $Bel(f(now) = x, s) \geq 1 - \epsilon$, for some value x and where s will be the current situation. Comparisons with the Bacchus *et al.* framework could also be made with more elaborate methods for combining noisy sensor data, for example Kalman filtering.

Bibliography

[1] Bacchus, Fahiem , Halpern, Joseph H. and Levesque, Hector J.(1998). Reasoning about noisy sensors and effectors in the situation calculus. *Artificial Intelligence*, vol. 111(1-2), pp. 171-208, 1999.

[2] Chellas, Brian F. Modal Logic. Cambridge University Press, Cambridge Massachusetts, 1980.

[3] del Val, Alvaro and Shoham, Yoav. A unified view of belief revision and update. *Journal of Logic and Computation*, vol. 4, pp. 797-810, 1994

[4] Engelson, S. and MacDermott, D. Error Correction in Mobile Robot Map Learning. Proceedings of the IEEE International Conference on Robotics and Automation (ICRA), 1992.

[5] Fox, D.,Burgard, W., and Thrun, S. Markov Localization for Mobile Robots in Dynamic Environments. *Journal of Artificial Intelligence Research*, vol. 11, pp. 391-427, 1999.

[6] Hintikka, J. Knowledge and Belief. Cornell University Press, Ithaca, NY., 1962.

[7] Hintikka, J. Semantics for propositional attitudes. Reference and Modality (Edited by Leonard Linsky), pp. 145-167. Oxford University Press, Oxford, United Kingdom, 1971.

[8] Kripke, S. A. Semantical analysis of modal logic. Zeitschrift fur Mathematische Logik und Grundlagen de Mathematik, vol. 9, pp. 67-96, 1963.

[9] Lenser, S. and Veloso, M. Sensor Resetting Localization for Poorly Modelled Mobile Robots. Proceedings of the IEEE International Conference on Robotics and Automation (ICRA), 2000.

[10] Levesque, Hector J. and Lakemeyer, Gerhard. The Logic of Knowledge Bases. The MIT Press, Massachusetts Institute of Technology, Cambridge, Massachusetts, 2000.

[11] McArthur, Gregory L. Reasoning about knowledge and belief: a survey. *Computational Intelligence*, vol. 4, pp. 223-243, 1988.

[12] McCarthy, John and Hayes, Patrick J. Some Philosophical Problems from the Standpoint of Artificial Intelligence. *Machine Intelligence*, vol. 4, pp. 463-502, 1969

[13] Maybeck, Peter S. Stochastic models, estimation, and control Volume 1. Academic Press, Inc., New York, New York, 1979.

[14] Moore, R.C. Reasoning about knowledge and action. Technical Note 191, SRI International, October 1980.

[15] Poole, David, Mackworth, Alan, and Goebel, Randy Computational Intelligence: a logical approach. Oxford University Press, Inc., Oxford, New York, 1998.

[16] Reiter, Raymond. Knowledge in Action: Logical Foundations for Specifying and Implementing Dynamical Systems. The MIT Press, Massachusetts Institute of Technology, Cambridge, Massachusetts, 2001.

[17] Reiter, Raymond. The Frame problem in the situation calculus: A simple solution (sometimes) and a completeness result for goal regression. Vladimir Lifschitz, editor. *Artificial Intelligence and Mathematical Theory of Computation: Papers in Honor of John McCarthy*, pp. 359-380. Academic Press, San Diego, CA, 1991.

[18] Russell, Stuart and Norvig, Peter Artificial Intelligence: A Modern Approach, 2nd Ed. Prentice Hall, Pearson Education, Inc., Upper Saddle River, New Jersey, 2003, 1995.

[19] Scherl, Richard B. and Levesque, Hector J. Knowledge, Action, and the Frame Problem *Artificial Intelligence*, vol. 144, pp. 1-39, 2003.

[20] Shapiro, Steven, Pagnucco, Maurice, Lespérance, Yves and Levesque, Hector J. Iterated Belief Change in the Situation Calculus.

[21] Thrun, S., Fox, D., Burgard, W. and Dellaert, F. Robust Monte Carlo Localization for Mobile Robots. *Artificial Intelligence*, vol. 128, 2001.

[22] Wolf, J., Burgard, W., and Burkhardt, H. Robust Vision-based Localization for Mobile Robots using an Image Retrieval System Based on Invariant Features. Proceedings of the IEEE International Conference on Robotics and Automation (ICRA), 2002.