
Reasoning about Probabilities in Dynamic Systems using Goal Regression*

Vaishak Belle and Hector J. Levesque
Dept. of Computer Science
University of Toronto
Toronto, Ontario M5S 3H5, Canada
{vaishak, hector}@cs.toronto.edu

Abstract

Reasoning about degrees of belief in uncertain dynamic worlds is fundamental to many applications, such as robotics and planning, where actions modify state properties and sensors provide measurements, both of which are prone to noise. With the exception of limited cases such as Gaussian processes over linear phenomena, belief state evolution can be complex and hard to reason with in a general way. This paper proposes a framework with new results that allows the reduction of subjective probabilities after sensing and acting, both in discrete and continuous domains, to questions about the initial state only. We build on an expressive probabilistic first-order logical account by Bacchus, Halpern and Levesque, resulting in a methodology that, in principle, can be coupled with a variety of existing inference solutions.

1 INTRODUCTION

Reasoning about *degrees of belief* in uncertain dynamic worlds is fundamental to many applications, such as robotics and planning, where actions modify state properties and sensors provide measurements, both of which are prone to noise. However, there seem to be two disparate paradigms to address this concern, both of which have their limitations. At one extreme, there are logical formalisms, such as the *situation calculus* (McCarthy and Hayes, 1969; Reiter, 2001), which allows us to express strict uncertainty, and exploits regularities in the effects actions have on propositions to describe physical laws compactly. Probabilistic sensor fusion, however, has received less attention here. At the other extreme, revising beliefs after noisy observations over rich error profiles is effortlessly addressed using *probabilistic techniques* such as Kalman filtering and Dynamic Bayesian Networks (Dean and Kanazawa, 1989; Dean and Wellman, 1991). However, in these frameworks, a complete specification of the

dependencies between variables is taken as given, making it difficult to deal with other forms of incomplete knowledge as well as complex actions that shift dependencies between variables in nontrivial ways.

An influential but nevertheless simple proposal by Bacchus, Halpern and Levesque (1999), BHL henceforth, was among the first to merge these broad areas in a general way. Their specification is widely applicable because it is not constrained to particular structural assumptions. In a nutshell, they extend the situation calculus language with a provision for specifying the degrees of belief in formulas in the initial state, closely fashioned after intuitions on incorporating probability in modal logics (Halpern, 1990; Fagin and Halpern, 1994). This then allows incomplete and partial specifications, which might be compatible with one or very many initial distributions and sets of independence assumptions, with beliefs following at a corresponding level of specificity. Moreover, together with a rich action theory, the model not only exhibits Bayesian conditioning (Pearl, 1988) (which, then, captures special cases such as Kalman filtering), but also allows flexibility in the ways dependencies and distributions may change over actions.

What is left open, however, is the following computational concern: how do we effectively reason about degrees of belief in the framework? That is, while changing degrees of belief do indeed emerge as *logical entailments* of the given action theory, no procedure is given for computing these entailments. On closer examination, in fact, this is a two-part question:

- (i) How do we effectively reason about beliefs in a particular state?
- (ii) How do we effectively reason about belief state evolution and belief change?

In the simplest case, part (i) puts aside acting and sensing, and considers reasoning about the *initial* state only, which is then the classical problem of (first-order) probabilistic inference. We do not attempt to do a full survey here, but this has received a lot of attention, often under reasonable assumptions such as the ability to factorize domains (Poole, 2003; Gogate and Domingos, 2010).

*We thank the reviewers for very helpful comments. Financial support from the Natural Science and Engineering Research Council of Canada made this research possible.

This paper is about part (ii). Addressing this concern has a critical bearing on the assumptions made about the domain for tractability purposes. For example, if the initial state supports a decomposed representation of the distribution, can we expect the same after actions? In the exception of very limited cases such as Kalman filtering that harness the conjugate property of Gaussian processes, the situation is discouraging. In fact, even in the slightly more general case of Dynamic Bayesian Networks, which are in essence atomic propositions, if one were to assume that state variables are independent at time 0, they can become fully correlated after a few steps (Dean and Kanazawa, 1988; Boyen and Koller, 1998; Hajishirzi and Amir, 2010). Dealing with complex actions, incomplete specifications and mixed representations, therefore, is significantly more involved.

In this paper, we propose a new alternative to infer degrees of belief in the presence of a rich theory of actions, closely related to *goal regression* (Waldinger, 1977; Reiter, 2001). The procedure is general, not requiring (but allowing) structural constraints about the domain, nor imposing (but allowing) limitations to the family of actions. Regression derives a mathematical formula, using term and formula *substitution* only, that relates belief after a sequence of actions and observations, even when they are noisy, to beliefs about the initial state. That is, among other things, if the initial state supports efficient factorizations, regression will maintain this advantage no matter how actions affect the dependencies between state variables over time. Going further, the formalism will work seamlessly with discrete probability distributions, probability densities, and perhaps most significantly, with difficult combinations of the two. (See Example 9.3 in Section 4 below.)

To see a simple example of what goal regression does, imagine a robot facing a wall and at a certain distance h to it, as in Figure 1. The robot might initially believe h to be drawn from a uniform distribution on $[2, 12]$. Assume the robot moves away by 2 units and is now interested in the belief about $h \leq 5$. Regression would tell the robot that this is equivalent to its initial beliefs about $h \leq 3$ which here would lead to a value of .1. To see a nontrivial example, imagine now the robot is also equipped with a sonar unit aimed at the wall, that adds Gaussian noise with mean μ and variance σ^2 . After moving away by 2 units, if the sonar were now to provide a reading of 8, then regression would derive that belief about $h \leq 5$ is equivalent to

$$\frac{1}{\gamma} \int_2^3 .1 \times \mathcal{N}(6 - x; \mu, \sigma^2) dx.$$

where γ is the normalization factor. Essentially, the posterior belief about $h \leq 5$ is reformulated as the product of the prior belief about $h \leq 3$ and the likelihood of $h \leq 3$ given an observation of 6. (That is, observing 8 after moving away by 2 units is related here to observing 6 initially.)

We believe the broader contributions of this line of work are

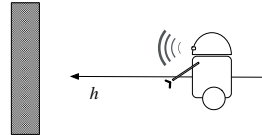


Figure 1: Robot moving towards a wall.

two-fold. On the one hand, as we show later, simple cases of belief state evolution, as applicable to conjugate distributions for example, are special cases of regression’s backward chaining procedure. Thus, regression could serve as a *formal basis* to study probabilistic belief change wrt limited forms of actions. On the other hand, our contribution can be viewed as a methodology for combining actions with recent advances in probabilistic inference, because reasoning about actions reduces to reasoning about the initial state.

We now describe the structure of the paper. Before moving on, we note that although the original BHL account is only suitable for discrete domains (as they assume values are taken from countable sets), in a companion paper (Belle and Levesque, 2013a) we show that the account can be generalized to domains with both discrete and continuous variables with minimal additions. In the preliminaries section, we cover the situation calculus, recap BHL and go over the essentials of its continuous extension. We then present regression for discrete domains, followed by regression for general domains. We end with related and future work. For this version of the paper, we allow noisy sensors but assume deterministic (noise-free) physical actions. Noisy actions are left for an extended version.

2 BACKGROUND

The language \mathcal{L} of the situation calculus (Reiter, 2001) is a many-sorted dialect of predicate calculus, with sorts for *physical actions*, *sensing actions*, *situations* and *objects* (including the set of reals \mathbb{R} as a subsort). A situation represents a history as a sequence of actions. A set of initial situations correspond to the ways the world might be initially. Successor situations are the result of doing actions, where the term $do(a, s)$ denotes the unique situation obtained on doing a in s . The term $do(\alpha, s)$, where α is the sequence $[a_1, \dots, a_n]$, abbreviates $do(a_n, do(\dots, do(a_1, s) \dots))$. For example, $do([grasp(o_1), repair(o_1)], s)$ represents the situation obtained after grasping and repairing object o_1 starting from s . Initial situations are those without a predecessor:

$$Init(s) \doteq \neg \exists a, s'. s = do(a, s').$$

We let the constant S_0 denote the actual initial situation, and we use the variable t to range over initial situations only. \mathcal{L} also includes functions whose values vary from situation to situation, called *fluents*, whose last argument is a situation.

We follow two notational conventions. We often suppress the situation argument in a formula ϕ , or use a distinguished

variable *now*. Either way, $\phi[t]$ is used to denote the formula with that variable replaced by t , e.g. both $(f < 12)[s]$ and $(f(now) < 12)[s]$ mean $f(s) < 12$. We also use conditional *if-then-else* expressions in formulas throughout. We write $f = \text{IF } \phi \text{ THEN } t_1 \text{ ELSE } t_2$ to mean $[\phi \wedge f = t_1] \vee [\neg\phi \wedge f = t_2]$. In case quantifiers appear inside the *if*-condition, we take some liberties with notation and the scope of variables in that we write $f = \text{IF } \exists x. \phi \text{ THEN } t_1 \text{ ELSE } t_2$ to mean $\exists x [\phi \wedge f = t_1] \vee [(f = t_2) \wedge \neg\exists x. \phi]$.

Basic action theory

Following (Reiter, 2001), we model dynamic domains in \mathcal{L} by means of a *basic action theory* \mathcal{D} , which consists of domain-independent *foundational* axioms, *unique name* axioms for actions (see (Reiter, 2001)), and (1) axioms \mathcal{D}_0 that describe what is true in the initial states, including S_0 ;¹ (2) precondition axioms² of the form $\text{Poss}(A(\vec{x}), s) \equiv \Pi_A(\vec{x}, s)$ describing executability conditions using a special fluent *Poss*; and (3) successor state axioms of the following form stipulating how fluents change:

$$f(\text{do}(a, s)) = u \equiv \exists \vec{z}_1 [a = A_1(\vec{z}_1) \wedge u = e_1(\vec{z}_1, s)] \vee \dots \\ \exists \vec{z}_k [a = A_k(\vec{z}_k) \wedge u = e_k(\vec{z}_k, s)] \vee \\ \wedge \neg \exists \vec{z}_i [a = A_i(\vec{z}_i)] \wedge u = f(s). \quad (1)$$

where $e_i(\vec{z}_i, s)$ is any expression whose only free variables are \vec{z}_i and s . For example, consider the action $\text{fwd}(z)$ of moving precisely z units towards or away from the wall, but the motion stops when the wall is reached:

$$h(\text{do}(a, s)) = u \equiv \exists z [a = \text{fwd}(z) \wedge u = \max(0, h(s) - z)] \vee \\ \neg \exists z [a = \text{fwd}(z)] \wedge u = h(s). \quad (2)$$

This sentence states that $\text{fwd}(z)$ is the only action affecting fluent h , in effect incorporating a solution to the frame problem (Reiter, 2001). Given an action theory, an agent reasons about actions by means of entailments of \mathcal{D} .³

Likelihood and belief

The BHL model of belief builds on a treatment of *knowledge* in \mathcal{L} (Scherl and Levesque, 2003). Here we present a simpler variant based on two distinguished fluents l and p .

The term $l(a, s)$ is intended to denote the likelihood of action a in situation s . For example, suppose $\text{sonar}(z)$ is the action of reading the value z from a sonar that measures the distance to the wall, h . We might assume that this action is characterized by a simple discrete error model (continuous

error models are considered later):

$$l(\text{sonar}(z), s) = \text{IF } |h(s) - z| \leq 1 \text{ THEN } 1/3 \text{ ELSE } 0 \quad (3)$$

which stipulates that the difference between a reading of z and the true value h is either $\{0, -1, 1\}$ with probability $1/3$, assuming that h and z take integer values. In general, the action theory \mathcal{D} is assumed to contain for each sensor $\text{sense}_i(\vec{x})$ that measures a fluent f , an axiom of the form:

$$l(\text{sense}_i(\vec{x}), s) = \text{Err}_i(\vec{x}, f(s)),$$

where $\text{Err}_i(u_1, u_2)$ is some expression with only two free variables u_1 and u_2 , both numeric.⁴ (Noise-free physical actions are given a likelihood of 1.)

Next, the p fluent determines a (subjective) probability distribution on situations. The term $p(s', s)$ denotes the relative *weight* accorded to situation s' when the agent happens to be in situation s , as in modal probability logics (Fagin and Halpern, 1994). Now, the task of the modeler is to specify the initial properties of p as part of \mathcal{D}_0 using ι and S_0 , e.g.:

$$p(\iota, S_0) = \text{IF } h(\iota) \in \{2, \dots, 11\} \text{ THEN } .1 \text{ ELSE } 0 \quad (4)$$

says that h is drawn from a uniform distribution. The following nonnegative constraint is also included in \mathcal{D}_0 :

$$\forall \iota, s. p(s, \iota) \geq 0 \wedge (p(s, \iota) > 0 \supset \text{Init}(s)) \quad (\text{P1})$$

Then, by means of a remarkably simple successor state axiom for p , (P2) below, the formal specification is complete.

$$p(s', \text{do}(a, s)) = \\ \text{IF } \exists s''. s' = \text{do}(a, s'') \wedge \text{Poss}(a, s'') \\ \text{THEN } p(s'', s) \times l(a, s'') \\ \text{ELSE } 0 \quad (\text{P2})$$

In particular, the *degree of belief* in a formula ϕ can be accounted for in terms of an abbreviation:⁵

$$\text{Bel}(\phi, s) \doteq \frac{1}{\gamma} \sum_{\{s': \phi[s']\}} p(s', s) \quad (\text{B})$$

where γ , the normalization factor, is understood throughout as the same expression as the numerator but with ϕ replaced by *true*, e.g. here γ is $\sum_{s'} p(s', s)$. So, as in probability logics, belief is simply the total weight of worlds satisfying ϕ . But the novelty here is that in a dynamical setting, belief change via (B) is identical to Bayesian conditioning:

¹Note that \mathcal{D}_0 can include any (classical) first-order sentence about S_0 , such as $h(S_0) > 12$ and $f_1(S_0) \neq 2 \vee f_2(S_0) = 5$.

²Free variables in any of these axioms should be understood as universally quantified from the outside.

³Entailments are wrt standard Tarskian models, but we will also assume that models assign the usual interpretations to $=, <, >, 0, 1, +, \times, /, -, e, \pi$, and x^y (exponentials).

⁴This captures the idea that the error model of a sensor measuring f depends only on the true value of f , and is independent of other factors. In a sense this follows the Bayesian model that conditioning on a random variable f is the same as conditioning on the event of observing f . But this is not required in general in the BHL scheme, an issue we ignore for this paper.

⁵Summations can be expressed as logical terms. See BHL.

Proposition 1: Suppose \mathcal{D} includes (P1), (P2) and the likelihood axiom for a sensor $sense(z)$ measuring f . Then

$$\mathcal{D} \models Bel(f = t, do(sense(z), S_0)) = \frac{Bel(f = t, S_0) \cdot Err(z, t)}{\sum_x Bel(f = x, S_0) \cdot Err(z, x)}$$

Essentially, if the robot’s sensors are informative, in the sense of returning values closer to the true value, beliefs are strengthened over time.

From sums to integrals

While the definition of belief in BHL has many desirable properties, it is defined in terms of a *summation* over situations, and therefore precludes fluents whose values range over the reals. The continuous analogue of (B) then requires *integrating* over some suitable space of values.

As it turns out, a suitable space can be found. First, assume that there are n fluents f_1, \dots, f_n in \mathcal{L} , and that these take no arguments other than the situation argument.⁶ Next, suppose that there is exactly one initial situation for every possible value of these fluents (Levesque *et al.*, 1998):

$$[\forall \vec{x} \exists \iota \bigwedge f_i(\iota) = x_i] \wedge [\forall \iota, \iota'. \bigwedge f_i(\iota) = f_i(\iota') \supset \iota = \iota'] \quad (\text{I})$$

Under these assumptions, it can be shown that the summation over all situations in (B) can be recast as a summation over all possible initial values x_1, \dots, x_n for the fluents:

$$Bel(\phi, s) \doteq \frac{1}{\gamma} \sum_{\vec{x}} P(\vec{x}, \phi, s) \quad (\text{B}')$$

where $P(\vec{x}, \phi, s)$ is the (unnormalized) weight accorded to the *successor* of an initial world where f_i equals x_i :

$$P(\vec{x}, \phi, do(\alpha, S_0)) \doteq \begin{array}{l} \text{IF } \exists \iota. \bigwedge f_i(\iota) = x_i \wedge \phi[do(\alpha, \iota)] \\ \text{THEN } p(do(\alpha, \iota), do(\alpha, S_0)) \\ \text{ELSE } 0 \end{array}$$

for any action sequence α . In a nutshell, because every situation has an initial situation as an ancestor, and because there is a bijection between initial situations and possible fluent values, it is sufficient to sum over fluent values to obtain the belief even for non-initial situations. Note that unlike (B), this one expects the final situation term $do(\alpha, S_0)$ mentioning what actions and observations took place to be explicitly specified, but that is just what one expects when the agent reasons about its belief after acting and sensing.

The generalization to the continuous case then proceeds as follows. First, we observe that some (though possibly not all) fluents will be real-valued, and that $p(s', s)$ will now be a measure of *density* not weight. For example, if h is real-valued, we might have the following analogue to (4):

$$p(\iota, S_0) = \text{IF } 2 \leq h(\iota) \leq 12 \text{ THEN } .1 \text{ ELSE } 0 \quad (5)$$

⁶It might be desirable to have fluents take arguments other than the situation. See (Belle and Levesque, 2013a) for discussions.

which says that the true initial value of h is drawn from a uniform distribution on $[2, 12]$. Similarly, the P term above now measures (unnormalized) density rather than weight.

Now suppose fluents are partitioned into two groups: the first k take their values x_1, \dots, x_k from \mathbb{R} , while the rest take their values y_{k+1}, \dots, y_n from countable domains, then the *degree of belief* in ϕ is an abbreviation for:

$$Bel(\phi, s) \doteq \frac{1}{\gamma} \int_{\vec{x}} \sum_y P(\vec{x} \cdot \vec{y}, \phi, s) \quad (\text{B}^+)$$

That is, the belief in ϕ is obtained by ranging over all possible fluent values, and integrating⁷ and summing the densities of *successor* situations where ϕ holds.⁸

To summarize the formalization, a basic action theory \mathcal{D} henceforth is assumed to additionally include: (a) (P1) and (I) as part of \mathcal{D}_0 ; (b) (P2) as part of \mathcal{D} ’s successor state axioms, and (c) sensor likelihood axioms.

3 REGRESSION FOR DISCRETE DOMAINS

We now investigate a computational mechanism for reasoning about beliefs after a trajectory. In this section, we focus on discrete domains, where a *weight*-based notion of belief would be appropriate. Domains with both discrete and continuous variables are reserved for the next section.

Formally, given a basic action theory \mathcal{D} , a sequence of actions α , we might want to determine whether a formula ϕ *holds* after executing α starting from S_0 :

$$\mathcal{D} \models \phi[do(\alpha, S_0)] \quad (6)$$

which is called *projection* (Reiter, 2001). When it comes to beliefs, and in particular how that changes after acting and sensing, we might be interested in *calculating* the degrees of belief in ϕ after α : find a real number n such that

$$\mathcal{D} \models Bel(\phi, do(\alpha, S_0)) = n. \quad (7)$$

The obvious method for answering (6) is to translate both \mathcal{D} and ϕ into a predicate logic formula. This approach, however, presents a serious computational problem because belief formulas expand into a large number of sentences using (P2), resulting in an enormous search space with initial and successor situations. The other issue with

⁷Like in BHL, where summations are captured as logical terms using second-order quantification, we can use logical formulas to capture a variety of sorts of integrals. See (Belle and Levesque, 2013a). We will henceforth simply suppose that for any term t and variable x , $\int_x t$ is a term which evaluates (in the standard calculus sense) to the integral of t between $[-\infty, \infty]$.

⁸We are assuming here that the density function is (Riemann) integrable. If it is not or if $\gamma = 0$ then belief is clearly not defined, nor should it be.

this approach is that sums (and integrals in the continuous case) reduce to complicated second-order formulas.

We now introduce a *regression* procedure to simplify both (6) and (7) to queries about $Bel(\phi, S_0)$, over arithmetic expressions, for which standard probabilistic reasoning methods can be applied. For this purpose, in the sequel, Bel is treated as a *special syntactic operator* rather than as an abbreviation for other formulas. To see a simple example of the procedure, imagine the robot is interested in the probability of $h=7$, given (4), after reading 5 from a sonar:

$$Bel(h = 7, do(sonar(5), S_0)) \quad (8)$$

If we are to take the sonar's model to be (3), then (8) should be 0 by Bayesian conditioning because the likelihood of the true value being 7 given an observation of 5 is 0. Regression would reduce the term (8) to one over initial priors:

$$\frac{1}{\gamma} \sum_{x \in \{2, \dots, 11\}} Err(5, x) \times Bel(h = x \wedge h = 7, S_0) \quad (9)$$

where Err is the error model from (3). By the condition inside Bel , the only valid value for x is 7 for which the prior is .1 but $Err(5, 7)$ is 0. Thus, (8) = (9) = 0. In general, regression is a *recursive* procedure that works iteratively over a sequence of actions discarding one action at a time, and it can be utilized to measure any logical property about the variables, e.g. $2\pi \cdot h < 12, h/fuel \leq \text{mileage}$, etc.

Formally, regression operates at two levels. (Note that this differs slightly from (Reiter, 2001; Scherl and Levesque, 2003).) At the formula level,⁹ we introduce an operator \mathcal{R} for regressing formulas, which over equality literals sends the individual terms to an operator \mathcal{T} for regressing terms. The fundamental objective of these operators is eliminate do symbols. The end result, then, is to transform any expression whose situation term is a successor of S_0 , say $do([a_1, a_2], S_0)$, to one about S_0 only, at which point \mathcal{D}_0 is all that is needed. As hinted earlier, these operators treat $Bel(\phi, s)$ as though they are special sorts of terms.¹⁰ Throughout the presentation, we assume that the inputs to these operators do not quantify over all situations.

Definition 2: For any term t , we inductively define $\mathcal{T}[t]$:

1. If t is situation-independent (e.g. $x, \pi^{2/3}$) then $\mathcal{T}[t] = t$.
2. $\mathcal{T}[g(t_1, \dots, t_k)] = g(\mathcal{T}[t_1], \dots, \mathcal{T}[t_k])$,
where g is any non-fluent function (e.g. $\times, +, \mathcal{N}$).

⁹For simplicity, in what follows, functional fluents in formulas are only allowed to occur as arguments of an equality literal. It is easy to show that every sentence can be transformed into an equivalent one in the required form, and the transformation is linear in the size of the original sentence, e.g. $h \leq 9$ is written as $\exists u (h = u \wedge u \leq 9)$.

¹⁰In the context of Bel , ϕ is understood to be any situation-suppressed formula not mentioning p, l and Bel . If situation terms do appear in ϕ , then they may only be the distinguished variable now .

3. For a fluent function f , $\mathcal{T}[f(s)]$ is defined inductively
 - (a) if s is of the form $do(A(\vec{t}), s')$ then
 $\mathcal{T}[f(s)] = \mathcal{T}[e(\vec{t}, s')]$
 - (b) else $\mathcal{T}[f(s)] = f(s)$

where, in (a), an appropriate instance of the *rhs* of the successor state axiom is used, as obtained from (1).

4. $\mathcal{T}[Bel(\phi, s)]$ is defined inductively:
 - (a) if s is of the form $do(a, s')$ and a is a physical action, then
 $\mathcal{T}[Bel(\phi, s)] = \mathcal{T}[Bel(\psi, s')]$
where ψ is $Poss(a, now) \supset \mathcal{R}[\phi(do(a, now))]$.
 - (b) if s is of the form $do(a, s')$ and a is a sensing action $sense(z)$ such that $l(sense(z), s) = Err(z, f_i(s))$ is in \mathcal{D} then
 $\mathcal{T}[Bel(\phi, s)] =$
 $\frac{1}{\gamma} \sum_{x_i} Err(z, x_i) \times \mathcal{T}[Bel(\psi, s')]$
where ψ is $Poss(a, now) \supset \phi \wedge f_i(now) = x_i$, and γ is the normalization factor and is the same expression as the numerator but ϕ replaced by $true$.
 - (c) else $\mathcal{T}[Bel(\phi, s)] = Bel(\phi, s)$.

Definition 3: For any formula ϕ , we define $\mathcal{R}[\phi]$ inductively:

1. $\mathcal{R}[t_1 = t_2] = (\mathcal{T}[t_1] = \mathcal{T}[t_2])$
2. $\mathcal{R}[G(t_1, \dots, t_k)] = G(\mathcal{T}[t_1], \dots, \mathcal{T}[t_k])$
where G is any non-fluent predicate (e.g. $=, <$).
3. $\mathcal{R}[Poss(A(\vec{t}), s)] = \mathcal{R}[\Pi_A(\vec{t}, s)]$,
where an appropriate instance of the *rhs* of the precondition axiom replaces the atom (see Section 2).
4. When ψ is a formula, $\mathcal{R}[\neg\psi] = \neg\mathcal{R}[\psi]$,
 $\mathcal{R}[\forall x\psi] = \forall x\mathcal{R}[\psi]$, $\mathcal{R}[\exists x\psi] = \exists x\mathcal{R}[\psi]$.
5. When ψ_1 and ψ_2 are formulas,
 $\mathcal{R}[\psi_1 \wedge \psi_2] = \mathcal{R}[\psi_1] \wedge \mathcal{R}[\psi_2]$,
 $\mathcal{R}[\psi_1 \vee \psi_2] = \mathcal{R}[\psi_1] \vee \mathcal{R}[\psi_2]$.

This completes the definition of \mathcal{T} and \mathcal{R} . We now go over the justifications for the items, starting with the operator \mathcal{T} . In item 1, non-fluents simply do not change after actions. In item 2, \mathcal{T} operates over sums and products in a modular manner. In item 3, provided there are remaining do symbols, the physics of the domain determines what the conditions must have been in the previous situation for the current value to hold. In item 4, if there is a remainder physical action, part (a) says that belief in ϕ after actions is simply the prior belief about the regression of ϕ , contingent on action executability. Part (b) says that the belief about

ϕ after observing z for the true value of f_i is the prior belief for all possible values x_i for f_i that agree with ϕ , times the likelihood of f_i being x_i given z . The appropriateness of parts (a) and (b) depend on the fact that physical actions do not have any sensing aspect, while sensing actions do not change the world. Part (c) simply says that \mathcal{T} stops when no *do* symbols appear in s . We proceed now with the justifications for \mathcal{R} . Over equality atoms, \mathcal{R} separates the terms of the equality and sends them to \mathcal{T} . Likewise, over non-fluent predicates. When *Poss* is encountered, preconditions take its place. Finally, \mathcal{R} simplifies over connectives in a straightforward way. The main result for \mathcal{R} regarding projection is:

Theorem 4: Suppose \mathcal{D} is any action theory, ϕ any situation-suppressed formula and α any action sequence:

$$\mathcal{D} \models \phi[do(\alpha, S_0)] \text{ iff } \mathcal{D}_0 \cup \mathcal{D}_{una} \models \mathcal{R}[\phi[do(\alpha, S_0)]]$$

where \mathcal{D}_{una} is the unique name assumption and $\mathcal{R}[\phi[do(\alpha, S_0)]]$ mentions only a single situation term, S_0 .

Here, \mathcal{D}_{una} is only needed to simplify action terms (Reiter, 2001) e.g. from $fwd(4) = fwd(z)$, \mathcal{D}_{una} infers $z = 4$. Now when our goal is to explicitly compute the degrees of belief in the sense of (7), we have the following property for \mathcal{T} :

Theorem 5: Let \mathcal{D} be as above, ϕ any situation-suppressed formula and α any sequence of actions. Then:

$$\mathcal{D} \models Bel(\phi, do(\alpha, S_0)) = \mathcal{T}[Bel(\phi, do(\alpha, S_0))]$$

where $\mathcal{T}[Bel(\phi, do(\alpha, S_0))]$ is a term about S_0 only.

Theorem 5 essentially shows how belief about trajectories is computable using beliefs about S_0 only. Note that, since the result of \mathcal{T} is a term about S_0 , no sentence outside of $\mathcal{D} - \mathcal{D}_0$ is needed. We now illustrate regression with examples. Using Theorem 5, we reduce beliefs after actions to initial ones. At the final step, standard probabilistic reasoning is applied to obtain the end values.

Example 6: Let \mathcal{D} contain the union of (2), (3) and (4).¹¹ Then the following equality expressions are entailed by \mathcal{D} :

$$1. Bel(h = 10 \vee h = 11, S_0) = .2$$

$$Bel(h \leq 9, S_0) = .8$$

Terms about S_0 are unaffected by \mathcal{T} . So this amounts to inferring probabilities using \mathcal{D}_0 .

$$2. Bel(h = 11, do(fwd(1), S_0))$$

$$= \mathcal{T}[Bel(h = 11, do(fwd(1), S_0))]$$

$$= \mathcal{T}[Bel(\mathcal{R}[(h = 11)][do(fwd(1), now)], S_0)] \quad (i)$$

¹¹Initial beliefs can also be specified for \mathcal{D}_0 using *Bel*, e.g. (4) can be replaced in \mathcal{D}_0 with $Bel(h = u, S_0) = .1$ for $u \in \{2, \dots, 11\}$.

$$= \mathcal{T}[Bel(\mathcal{T}[h(do(fwd(1), now))]) = \mathcal{T}[11], S_0)] \quad (ii)$$

$$= \mathcal{T}[Bel(\max(0, h - 1) = 11, S_0)] \quad (iii)$$

$$= Bel(\max(0, h - 1) = 11, S_0) \quad (iv)$$

$$= 0$$

First, since action preconditions are all true, *Poss* is ignored everywhere. We underline to emphasize the expressions undergoing transformations. We begin always by applying \mathcal{T} to the main term, in this case getting (i), by means of \mathcal{T} 's item 4(a). Next, \mathcal{R} 's item 1 is applied in (ii). While $\mathcal{T}[11] = 11$ by \mathcal{T} 's item 1, for $\mathcal{T}[h(do(fwd(1), now))]$ we use item 3 and (2) to get:

$$\mathcal{T}[\max(0, h(now) - 1)] = \max(0, h(now) - 1)$$

which is substituted in (ii) to give (iii). Finally, \mathcal{T} 's item 4(c) yields (iv), which is a belief term about S_0 . Now the only valid value for h in (iv) is 12, but for $h = 12$ the robot has a belief of 0 initially.

$$3. Bel(h \leq 5, do(sonar(5), S_0))$$

$$= \frac{1}{\gamma} \sum_{x \in \{2, \dots, 11\}} Err(5, x) \times \mathcal{T}[Bel(h = x \wedge h \leq 5, S_0)] \quad (i)$$

$$= \frac{1}{\gamma} \sum_{x \in \{2, \dots, 11\}} Err(5, x) \times Bel(h = x \wedge h \leq 5, S_0) \quad (ii)$$

$$= \frac{1}{\gamma} \left(\frac{1}{3} \cdot Bel(h = 4 \wedge h \leq 5, S_0) + \frac{1}{3} \cdot Bel(h = 5 \wedge h \leq 5, S_0) + \frac{1}{3} \cdot Bel(h = 6 \wedge h \leq 5, S_0) \right) \quad (iii)$$

$$= \frac{1}{\gamma} \left(\frac{1}{3} \cdot Bel(h = 4, S_0) + \frac{1}{3} \cdot Bel(h = 5, S_0) \right) \quad (iv)$$

$$= \frac{1}{\gamma} \cdot \frac{2}{30}$$

$$= 2/3$$

where $Err(5, x)$ is the model from (3). First, \mathcal{T} 's item 4(b) yields (i), and then item 4(c) yields (ii). Since $Err(5, x)$ is non-zero only for $x \in \{4, 5, 6\}$, (ii) is simplified to (iii) and (iv) resulting in $1/15 \cdot 1/\gamma$. We calculate γ as follows:

$$= \sum_{x \in \{2, \dots, 11\}} Err(5, x) \times \mathcal{T}[Bel(h = x \wedge true, S_0)] \quad (i')$$

$$= \sum_{x \in \{2, \dots, 11\}} Err(5, x) \times Bel(h = x, S_0) \quad (ii')$$

$$= 3/30.$$

4 REGRESSION FOR GENERAL DOMAINS

We now generalize regression for domains with discrete and continuous variables, for which a *density*-based notion

of belief is appropriate. The main issue is that when formulating posterior beliefs after sensing, something like Definition 2's item 4(b) will not work. This is because over continuous spaces the belief about any individual point is 0. Therefore, we will be unpacking belief in terms of the density function, *i.e.* in terms of P . These $P(\vec{x}, \phi, s)$ terms, which will now also be treated as special sorts of syntactic terms, are separately regressed. (Of course, the regression of weight-based belief can be approached on similar lines.) Recall that $P(\vec{x}, \phi, S_0)$ is simply the *density* of an initial world (where $f_i = x_i$) satisfying ϕ . Formally, term regression \mathcal{T} is defined as follows:

Definition 7: For any term t , we inductively define:

1, 2 and 3 as before.

4. $\mathcal{T}[P(\vec{x}, \phi, s)]$ is defined inductively:

(a) if s is of the form $do(a, s')$ and a is a physical action then

$$\mathcal{T}[P(\vec{x}, \phi, s)] = \mathcal{T}[P(\vec{x}, \psi, s')]$$

where ψ is $Poss(a, now) \supset \mathcal{R}[\phi[do(a, now)]]$.

(b) if s is of the form $do(a, s')$ and a is a sensing action $sense(z)$ such that $l(sense(z), s) = Err(z, f_i(s))$ is in \mathcal{D} , then:

$$\mathcal{T}[P(\vec{x}, \phi, s)] = Err(z, x_i) \times \mathcal{T}[P(\vec{x}, \psi, s')]$$

where ψ is $Poss(a, now) \supset \phi \wedge f_i(now) = x_i$.

(c) else $\mathcal{T}[P(\vec{x}, \phi, s)] = P(\vec{x}, \phi, s)$.

$$5. \mathcal{T}[Bel(\phi, s)] = \frac{1}{\gamma} \int_{\vec{z}} \sum_y \mathcal{T}[P(\vec{z}, \vec{y}, \phi, s)].$$

\mathcal{R} is defined as before. We have the following property:

Theorem 8: Let \mathcal{D} be any action theory, ϕ any situation-suppressed formula and α any action sequence. Then

$$\mathcal{D} \models Bel(\phi, do(\alpha, S_0)) = \mathcal{T}[Bel(\phi, do(\alpha, S_0))]$$

where $\mathcal{T}[Bel(\phi, do(\alpha, S_0))]$ is a term about S_0 only.

Similarly, the analogue of Theorem 4 holds as well.

Example 9: Consider the following *continuous* variant of the robot example. Imagine a continuous uniform distribution for the true value of h , as provided by (5). Suppose the sonar has the following error profile:

$$l(sonar(z), s) = \begin{array}{l} \text{IF } z \geq 0 \\ \quad \text{THEN } \mathcal{N}(z - h(s); 0, 4) \\ \quad \text{ELSE } 0 \end{array} \quad (10)$$

which says the difference between a nonnegative reading and the true value is normally distributed with mean 0 and variance 4. (A mean of 0 implies there is no systematic bias.) Now, let \mathcal{D} be any action theory that includes (2), (5) and (10). Then the following equalities are entailed by \mathcal{D} :

$$1. Bel(h = 3 \vee h = 4, S_0) = 0,$$

$$Bel(4 \leq h \leq 6, S_0) = .2$$

\mathcal{T} does not change terms about S_0 . Here, for example, the second belief term equals $\int_4^6 .1 dx = .2$.

$$2. Bel(h \geq 11, do(fwd(1), S_0))$$

$$= \frac{1}{\gamma} \int_{x \in \mathbb{R}} \mathcal{T}[P(x, \underline{h \geq 11}, do(fwd(1), S_0))] \quad (i)$$

$$= \frac{1}{\gamma} \int_{x \in \mathbb{R}} \mathcal{T}[P(x, \mathcal{R}[\underline{\psi}], S_0)] \quad (ii)$$

where ψ is $(h \geq 11)[do(fwd(1), now)]$

$$= \frac{1}{\gamma} \int_{x \in \mathbb{R}} \mathcal{T}[P(x, \max(0, h - 1) \geq 11, S_0)] \quad (iii)$$

$$= \frac{1}{\gamma} \int_{x \in \mathbb{R}} P(x, \max(0, h - 1) \geq 11, S_0) \quad (iv)$$

$$= \frac{1}{\gamma} \int_{x \in \mathbb{R}} \begin{cases} p(\iota, S_0) & \text{if } \exists \iota. h(\iota) = x \wedge h(\iota) \geq 12 \\ 0 & \text{otherwise} \end{cases} \quad (v)$$

$$= \frac{1}{\gamma} \int_{x \in \mathbb{R}} \begin{cases} .1 & \text{if } x \in [2, 12] \text{ and } x \geq 12 \\ 0 & \text{otherwise} \end{cases} \quad (vi)$$

$$= \frac{1}{\gamma} \int_{x \in \mathbb{R}} \begin{cases} .1 & \text{if } x = 12 \\ 0 & \text{otherwise} \end{cases} \quad (vii)$$

$$= 0$$

We use \mathcal{T} 's item 5 to get (i), after which item 4(a) is applied. On doing \mathcal{R} in (ii), along the lines of Example 6.2, we obtain (iii). \mathcal{T} 's item 4(c) then yields (iv), and stops. In the steps following (iv), we show how P expands in terms of p , and how the space of situations resolves into a mathematical expression, yielding 0.¹²

$$3. Bel(h = 0, do(fwd(4), S_0))$$

$$= \frac{1}{\gamma} \int_{x \in \mathbb{R}} \mathcal{T}[P(x, \mathcal{R}[(h = 0)[do(fwd(4), now)]], S_0)] \quad (i)$$

$$= \frac{1}{\gamma} \int_{x \in \mathbb{R}} \mathcal{T}[P(x, \max(0, h - 4) = 0, S_0)] \quad (ii)$$

$$= \frac{1}{\gamma} \int_{x \in \mathbb{R}} \begin{cases} .1 & \text{if } x \in [2, 12] \text{ and } x \leq 4 \\ 0 & \text{otherwise} \end{cases} \quad (iii)$$

$$= .2$$

By means of (2), after moving forward by 4 units the belief about h is characterized by a *mixed* distribution because $h = 0$ is accorded a .2 weight (*i.e.* from all points where $h \in [2, 4]$ initially), while $h \in (0, 8]$ are associated with a density of .1. Here, \mathcal{T} 's item 5 and 4(a) are triggered, and the removal of \mathcal{T} using 4(c) is not shown. The end result is that the density function is integrated for $2 \leq x \leq 4$ leading to .2. (γ is 1.)

¹²Given certain assumptions, it is possible to further reduce logical expressions involving fluents to a mathematical expression using only those variables that appear in the integral. We expand on this in a longer version of the paper.

$$\begin{aligned}
4. & \text{Bel}(h = 4, do(fwd(-4), do(fwd(4), S_0))) \\
&= \frac{1}{\gamma} \int_{x \in \mathbb{R}} \mathcal{T}[P(x, \exists u. h = u \wedge \\
&\quad 4 = \max(0, u + 4), do(fwd(4), S_0))] \quad (i) \\
&= \frac{1}{\gamma} \int_{x \in \mathbb{R}} \mathcal{T}[P(x, \exists u. u = \max(0, h - 4) \wedge \\
&\quad 4 = \max(0, u + 4), S_0)] \quad (ii) \\
&= \frac{1}{\gamma} \int_{x \in \mathbb{R}} \begin{cases} .1 & \text{if } x \in [2, 12], x \leq 4 \\ 0 & \text{otherwise} \end{cases} \quad (iii) \\
&= .2
\end{aligned}$$

We noted above that the point $h = 4$ gets a .2 weight on executing $fwd(4)$, after which it obtains a h value of 0. The weight is *retained* on reversing by 4 units, with the point now obtaining a h value of 4. The derivation invokes two applications of \mathcal{T} 's item 4(a). We skip the intermediate \mathcal{R} steps. (γ evaluates to 1.)

$$\begin{aligned}
5. & \text{Bel}(h = 4, do(fwd(4), do(fwd(-4), S_0))) \\
&= \frac{1}{\gamma} \int_{x \in \mathbb{R}} \mathcal{T}[P(x, \exists u. u = \max(0, h + 4) \wedge \\
&\quad 4 = \max(0, u - 4), S_0)] \quad (i) \\
&= 0
\end{aligned}$$

Had the robot moved away first, no ‘‘collapsing’’ of points takes place, h remains a continuous distribution and no point is accorded a non-zero weight. \mathcal{T} steps are skipped but they are symmetric to the one above, *e.g.* compare (i) here and (ii) above. But then the density function is non-zero only for the individual $h = 4$.

$$\begin{aligned}
6. & \text{Bel}(4 \leq h \leq 6, do(sonar(5), S_0)) \\
&= \frac{1}{\gamma} \int_{x \in \mathbb{R}} \mathcal{N}(5 - x; 0, 4) \times \mathcal{T}[P(x, \psi, S_0)] \quad (i) \\
&\quad \text{where } \psi \text{ is } h = x \wedge 4 \leq h \leq 6 \\
&= \frac{1}{\gamma} \int_{x \in \mathbb{R}} \begin{cases} .1 \cdot \mathcal{N}(5 - x; 0, 4) & \text{if } x \in [2, 12], x \in [4, 6] \\ 0 & \text{otherwise} \end{cases} \\
&\approx .41
\end{aligned}$$

We obtain (i) after \mathcal{T} 's item 5 and then 4(b) for sensing actions. That is, belief about $h \in [4, 6]$ is sharpened after observing 5. Basically, we are integrating a function that is 0 everywhere except when $4 \leq x \leq 6$ where it is $.1 \times \mathcal{N}(5 - x; 0, 4)$, normalized over $2 \leq x \leq 12$.

$$\begin{aligned}
7. & \text{Bel}(4 \leq h \leq 6, do(sonar(5), do(sonar(5), S_0))) \\
&= \frac{1}{\gamma} \int_{x \in \mathbb{R}} \mathcal{N}(5 - x; 0, 4) \times \mathcal{T}[P(x, \psi, s)] \quad (i) \\
&\quad \text{where } s = do(sonar(5), S_0), \psi \text{ is } h = x \wedge 4 \leq h \leq 6 \\
&= \frac{1}{\gamma} \int_{x \in \mathbb{R}} [\mathcal{N}(5 - x; 0, 4)]^2 \times \mathcal{T}[P(x, \psi, S_0)] \quad (ii) \\
&\approx .52
\end{aligned}$$

As expected, two successive observations of 5 sharpens belief further. Derivations (i) and (ii) follow from

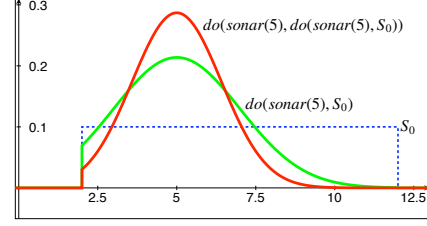


Figure 2: Belief density change for h at S_0 (in blue), after sensing 5 (in green) and after sensing 5 twice (in red).

\mathcal{T} 's item 5, and two successive applications of item 4(b). Thus, we are to integrate $.1 \times [\mathcal{N}(5 - x; 0, 4)]^2$ between $[4, 6]$ and normalize over $[2, 12]$. These changing densities are plotted in Figure 2.

5 TWO SPECIAL CASES

Regression is a general property for computing properties about posteriors in terms of priors after actions. It is therefore possible to explore limited cases, which might be appropriate for some applications. We present two such cases.

Conjugate distributions

Certain types of systems, such as Gaussian processes, admit an effective propagation model. The same advantages can be observed in our framework. We illustrate this using an example. Assume a fluent f , and suppose \mathcal{D}_0 is the union of (I), (P1) and the following specification:

$$p(t, S_0) = \mathcal{N}(f(t); \mu_1, \sigma_1^2)$$

which stipulates that the true value of f is believed to be normally distributed. Assume the following sensor in \mathcal{D} :

$$l(\text{sense}(z), s) = \mathcal{N}(z - f(s); \mu_2, \sigma_2^2)$$

Then it is easy to show that estimating posteriors yields a product of Gaussians (that is also a Gaussian process (Box and Tiao, 1973)), which is inferred by \mathcal{T} :¹³

$$\begin{aligned}
&\mathcal{T}[\text{Bel}(b \leq f \leq c, do(\text{sense}(z), S_0))] = \\
&\quad \frac{1}{\gamma} \int_b^c \mathcal{N}(x; \mu_1, \sigma_1^2) \cdot \mathcal{N}(z - x; \mu_2, \sigma_2^2) dx
\end{aligned}$$

Distribution transformations

Certain actions affect priors in a characteristically simple manner, and regression would account for these changes as an appropriate function of the initial belief state. We illustrate two instances using Example 9. First, consider an

¹³This corresponds to a simple case of Kalman filtering (Dean and Wellman, 1991), where the sensed value is static. In the complete framework with noisy effectors, we would obtain a model where distinct actions may condition priors in distinct ways.

action $grasp(z)$ that grabs object z . Because the action of grasping does not affect h by way of (2), we get:

$$\mathcal{T}[Bel(h \leq b, do(grasp(obj5), S_0))] = Bel(h \leq b, S_0)$$

So no changes to h 's density are required. Second, consider ground actions with the property that two distinct values of f do not become the same after that action, *e.g.*, for initial states this means:

$$\forall t, t'. f(t) \neq f(t') \supset f(do(a, t)) \neq f(do(a, t')) \quad (\text{EQ})$$

Think of $fwd(-4)$ that agrees with this, but $fwd(4)$ need not. We can show that such actions “shift” priors:

$$\mathcal{T}[Bel(h \leq b, do(fwd(-n), S_0))] = Bel(h \leq b - n, S_0)$$

Intuitively, the probability of h being in the interval $[b, c]$, irrespective of the distribution family, is the same as the probability of $h \in [b + n, c + n]$ after $fwd(-n)$. Thus, regression derives the initial interval given the current one.¹⁴

6 RELATED WORK

Perhaps the most popular models to treat sensor fusion include variants of Kalman filtering (Fox *et al.*, 2003; Thrun *et al.*, 2005), where priors and likelihoods are assumed to be Gaussian. We already pointed out some instances of Kalman filtering in our example. Where we differ is that backward chaining is possible even when: (a) no assumptions about the nature of distributions, nor about how distributions and dependencies change need to be made, (b) the framework is embedded in a rich theory of actions, and (c) arbitrary forms of incomplete knowledge are allowed, including strict uncertainty.¹⁵ Domain-specific dependencies, then, may be exploited as appropriate.

There have been, of course, other attempts to extend the situation calculus to reason about probabilistic belief, such as (Poole, 1998). See BHL for a discussion on the differences to that work. On a related note, there are numerous approaches that combine logic and probability. In particular, we mention dynamic logic proposals (Van Benthem *et al.*, 2009), planning languages (Younes and Littman, 2004; Sanner, 2011; Kushmerick *et al.*, 1995), and first-order frameworks based on the situation calculus and close relatives (Thielscher, 2001). For discussions on these and non-dynamic proposals such as Markov logics (Richardson and Domingos, 2006), see (Belle and Levesque, 2013a,b).

¹⁴Although fwd shifts the distribution linearly, in general, provided something like (EQ) is true, actions may also result in nonlinear changes to state variables, which would nonlinearly change the mean of the distribution. Nevertheless, similar features would be observed. See the longer version of the paper for more details.

¹⁵Moreover, the BHL model is compatible with a wide variety of formalisms such as (Fagin and Halpern, 1994; Halpern and Tuttle, 1993; Halpern, 1990). See BHL for discussions.

There is one other thread of related work, that of *symbolic dynamic programming* (Boutilier *et al.*, 2000, 2001) which also has recent continuous extensions (Sanner *et al.*, 2011). While regression is used in this literature as well, the concerns are very different: they focus on policy generation, while ours is strictly about *belief change*. Consequently, the regression in that literature is adapted from the regression for the non-epistemic situation calculus (Reiter, 2001). Ours, on the other hand, continues in the tradition of the epistemic situation calculus (Scherl and Levesque, 2003) by extending those intuitions to probabilistic belief and noisy sensing. In this regard, our account allows the modeler to explicitly reason about beliefs in the language, which would prove useful in formalizing the achievability of plans (Levesque, 1996), among other things. The idea of regression is not new and lies at the heart of many planning systems (Fritz and McIlraith, 2007). For STRIPS actions, regression has at most linear complexity in the length of the action sequence (Reiter, 2001). For other studies, see (Van Ditmarsch *et al.*, 2007; Rintanen, 2008).

7 CONCLUSIONS

Planning and robotic applications have to deal with numerous sources of complexity regarding action and change. Consequently, irrespective of the decompositions and factorizations that are justifiable initially, belief state evolution is known to invalidate these efforts even over simple temporal phenomena. In this paper, we obtain a general methodology to relate beliefs after acting and sensing to initial beliefs. We investigated the methodology in an existing model by BHL, and a continuous extension to it, making the technique applicable to discrete domains as well as general ones. We demonstrated regression using an example where actions affect priors in nonstandard ways, such as transforming a continuous distribution to a mixed one. In general, regression does not insist on (but allows) restrictions to actions, that is, no assumptions need to be made about how actions affect variables and their dependencies over time. Moreover, at the specification level, we do not assume (but allow) structurally constrained initial states.

There are many avenues for future work. Extending automated regression solutions (Reiter, 2001) to subjective probabilities is ongoing work. Moreover, given the promising advances made in the area of relational probabilistic inference, we believe regression suggests natural ways to apply those developments with actions. This line of research would allow us to address effective belief propagation for numerous planning problems that require both logical and probabilistic representations. On another front, note that after applying the reductions, one may also use approximate inference methods. Perhaps then, regression can serve as a computational framework to study approximate belief propagation, on the one hand, and using approximate inference at the initial state after goal regression, on the other.

References

- F. Bacchus, J. Y. Halpern, and H. J. Levesque. Reasoning about noisy sensors and effectors in the situation calculus. *Artificial Intelligence*, 111(1–2):171–208, 1999.
- V. Belle and H. J. Levesque. Reasoning about continuous uncertainty in the situation calculus. In *Proc. IJCAI*, 2013.
- V. Belle and H. J. Levesque. Robot location estimation in the situation calculus. In *ICAPS Workshop on Planning and Robotics*. 2013.
- C. Boutilier, R. Reiter, M. Soutchanski, and S. Thrun. Decision-theoretic, high-level agent programming in the situation calculus. In *Proc. AAAI*, pages 355–362, 2000.
- C. Boutilier, R. Reiter, and B. Price. Symbolic dynamic programming for first-order MDPs. In *Proc. IJCAI*, pages 690–697, 2001.
- G. E. P. Box and G. C. Tiao. *Bayesian inference in statistical analysis*. Addison-Wesley, 1973.
- X. Boyen and D. Koller. Tractable inference for complex stochastic processes. In *Proc. UAI*, pages 33–42, 1998.
- T. Dean and K. Kanazawa. Probabilistic temporal reasoning. In *Proc. AAAI*, pages 524–529, 1988.
- T. Dean and K. Kanazawa. A model for reasoning about persistence and causation. *Computational intelligence*, 5(2):142–150, 1989.
- T. Dean and M. Wellman. *Planning and control*. Morgan Kaufmann Publishers Inc., 1991.
- R. Fagin and J. Y. Halpern. Reasoning about knowledge and probability. *J. ACM*, 41(2):340–367, 1994.
- D. Fox, J. Hightower, L. Liao, D. Schulz, and G. Borriello. Bayesian filtering for location estimation. *Pervasive Computing, IEEE*, 2(3):24–33, 2003.
- C. Fritz and S. A. McIlraith. Monitoring plan optimality during execution. In *Proc. ICAPS*, pages 144–151, 2007.
- V. Gogate and P. Domingos. Formula-based probabilistic inference. In *Proc. UAI*, pages 210–219, 2010.
- H. Hajishirzi and E. Amir. Reasoning about deterministic actions with probabilistic prior and application to stochastic filtering. In *Proc. KR*, 2010.
- J. Y. Halpern and M. R. Tuttle. Knowledge, probability, and adversaries. *J. ACM*, 40:917–960, 1993.
- J.Y. Halpern. An analysis of first-order logics of probability. *Artificial Intelligence*, 46(3):311–350, 1990.
- N. Kushmerick, S. Hanks, and D.S. Weld. An algorithm for probabilistic planning. *Artificial Intelligence*, 76(1):239–286, 1995.
- H. J. Levesque, F. Pirri, and R. Reiter. Foundations for the situation calculus. *Electron. Trans. Artif. Intell.*, 2:159–178, 1998.
- H. J. Levesque. What is planning in the presence of sensing? In *Proc. AAAI/IAAI*, pages 1139–1146, 1996.
- J. McCarthy and P. J. Hayes. Some philosophical problems from the standpoint of artificial intelligence. In *Machine Intelligence*, pages 463–502, 1969.
- J. Pearl. *Probabilistic reasoning in intelligent systems: networks of plausible inference*. Morgan Kaufmann, 1988.
- D. Poole. Decision theory, the situation calculus and conditional plans. *Electron. Trans. Artif. Intell.*, 2:105–158, 1998.
- D. Poole. First-order probabilistic inference. In *Proc. IJCAI*, pages 985–991, 2003.
- R. Reiter. *Knowledge in action: logical foundations for specifying and implementing dynamical systems*. MIT Press, 2001.
- M. Richardson and P. Domingos. Markov logic networks. *Machine learning*, 62(1):107–136, 2006.
- J. Rintanen. Regression for classical and nondeterministic planning. In *Proc. ECAI*, pages 568–572, 2008.
- S. Sanner, K. V. Delgado, and L. N. de Barros. Symbolic dynamic programming for discrete and continuous state MDPs. In *Proc. UAI*, pages 643–652, 2011.
- S. Sanner. Relational dynamic influence diagram language (rddl): Language description. Technical report, Australian National University, 2011.
- R. B. Scherl and H. J. Levesque. Knowledge, action, and the frame problem. *Artificial Intelligence*, 144(1-2):1–39, 2003.
- M. Thielscher. Planning with noisy actions (preliminary report). In *Proc. Australian Joint Conference on Artificial Intelligence*, pages 27–45, 2001.
- S. Thrun, W. Burgard, and D. Fox. *Probabilistic Robotics*. MIT Press, 2005.
- J. Van Benthem, J. Gerbrandy, and B. Kooi. Dynamic update with probabilities. *Studia Logica*, 93(1):67–96, 2009.
- H. Van Ditmarsch, A. Herzig, and T. De Lima. Optimal regression for reasoning about knowledge and actions. In *Proc. AAAI*, pages 1070–1075, 2007.
- R. Waldinger. Achieving several goals simultaneously. In *Machine Intelligence*, volume 8, pages 94–136. 1977.
- H. Younes and M. Littman. PPDDL 1. 0: An extension to pddl for expressing planning domains with probabilistic effects. Technical report, Carnegie Mellon University, 2004.