Journal of Applied Logic 13 (2015) 397–413

Contents lists available at ScienceDirect

Journal of Applied Logic

www.elsevier.com/locate/jal

Robot location estimation in the situation calculus ☆

Vaishak Belle *, Hector J. Levesque

Dept. of Computer Science, University of Toronto, Toronto, Ontario M5S 3H5, Canada

A R T I C L E I N F O

Article history:
Received 19 February 2014
Available online 23 February 2015

Keywords:
Knowledge representation
Cognitive robotics
Reasoning about knowledge and uncertainty

A B S T R A C T

Location estimation is a fundamental sensing task in robotic applications, where the world is uncertain, and sensors and effectors are noisy. Most systems make various assumptions about the dependencies between state variables, and especially about how these dependencies change as a result of actions. Building on a general framework by Bacchus, Halpern and Levesque for reasoning about degrees of belief in the situation calculus, and a recent extension to it for continuous probability distributions, in this paper we illustrate location estimation in the presence of a rich theory of actions using examples. The formalism also allows specifications with incomplete knowledge and strict uncertainty, as a result of which the agent's initial beliefs need not be characterized by a unique probability distribution. Finally, we show that while actions might affect prior distributions in nonstandard ways, suitable posterior beliefs are nonetheless entailed as a side-effect of the overall specification.

1. Introduction

Location estimation is a fundamental sensing task in robotic applications [53], where the robot has to situate itself in an incompletely known world. To act purposefully in this setting, agents grapple with at least two sorts of reasoning problems. First, because the world is *dynamic*, actions perpetually change the properties of the state. Second, because little in the world is definite, the agent has to modify its beliefs based on the actions performed and the measurements on its sensors, both of which are prone to noise.

To see an illustrative example, imagine a robot operating in a 2-dimensional world, and located at a certain distance h to a wall, as in Fig. 1. Suppose the robot initially believes that h is drawn from a uniform distribution on the interval $[2, 12]$. Among the robot's many capabilities, suppose the robot can move along the axes. A leftward motion of one unit would shift the uniform distribution on h to $[1, 11]$, but a leftward motion of five units would change the distribution more radically. The point $h = 0$ would obtain a weight

☆ Preliminary versions of this work appear in *Proceedings of the Symposium on Logical Formalizations on Commonsense Reasoning*, 2013, and in *Proceedings of the Workshop on Weighted Logics for Artificial Intelligence*, 2013.

* Corresponding author.
E-mail addresses: vaishak@cs.toronto.edu (V. Belle), hector@cs.toronto.edu (H.J. Levesque).

http://dx.doi.org/10.1016/j.jal.2015.02.004
1570-8683/

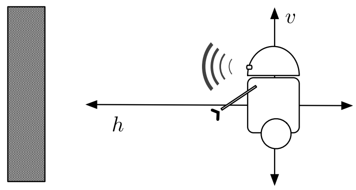

Fig. 1. Robot operating in a 2-dimensional world.

of .3, while $h \in (0, 7]$ would retain their densities. This mixed distribution would then be preserved by a subsequent rightward motion.

We might imagine, likewise, that the robot is equipped with two onboard sensors: a laser and a sonar unit, both aimed at the wall and so estimating h. Each of these might be characterized by Gaussian error models, and the effect of a reading from any sensor would revise the distribution on h from uniform to an appropriate Gaussian. In the end, the robot is left with the difficult task of adjusting its beliefs as it operates and obtains competing (perhaps conflicting) measurements from individual sensors.

Of course, if we were to assume instead that the robot believes its distance to the wall is uniformly distributed on the interval [2, 12] or on the interval [10, 20] without being able to say which, then we would still expect the robot to function as usual. The observation, then, is that the robot's belief about any formula should take both these distributions into account.

Conventional methods in robotics, such as the ones treated in [53], provide powerful algorithmic machinery in service of location estimation and navigation. Typically, in these models, one assumes that the agent's initial beliefs about the world are characterized by random variables drawn from well-known distributions, and then the dependencies between these variables are given by the domain modeler. For such specifications, revising beliefs after noisy observations and similar temporal phenomena is addressed using probabilistic techniques such as Kalman filtering and Dynamic Bayesian Networks [16,17]. But while belief update of known priors over Gaussian and other continuous error models is treated appropriately there, very little is said about how actions might change values of certain state variables while not affecting others. In particular, simplistic assumptions are often made about how dependencies change as a result of actions. Moreover, many actions in the real world often depend on complicated features of the environment – for example, a sensor's noise model might depend on the temperature of the room – which is cumbersome to handle in these models. Most significantly, since these formalisms assume a full specification of the dependencies between variables, it is difficult to deal with other forms of incomplete knowledge, like strict uncertainty (as can be obtained by disjunctions, for example). So, even for the simple setting in Fig. 1, things like shifting densities (such as the changing nature of the distribution on h), shifting dependencies (such as a temperature-dependent sensor model), and disjunctive knowledge cannot be captured easily using these probabilistic formalisms.

Perhaps the most general formalism for dealing with probabilistic beliefs in formulas, and how that should evolve in the presence of noisy acting and sensing, is a logical account by Bacchus, Halpern and Levesque (BHL) [2]. In the BHL approach, besides quantifiers and other logical connectives, one has the provision for specifying the *degrees of belief* in formulas in the initial state. This specification may be compatible with one or very many initial distributions and sets of independence assumptions. All the properties of belief will then follow at a corresponding level of specificity.

Subjective uncertainty is captured in the BHL approach using a possible-world model of belief [30,25,19]. Intuitively, the degree of belief in ϕ is defined as a normalized sum over the possible worlds where ϕ is true of some nonnegative *weights* associated with those worlds. To reason about belief change, the BHL model is then embedded in a rich theory of action and sensing provided by the situation calculus [37,46,50]. The BHL account provides axioms in the situation calculus regarding how the weight associated with a possible world changes as the result of acting and sensing. The properties of belief and belief change then emerge as a direct logical consequence of the initial specifications and these changes in weights.