Artificial Intelligence 262 (2018) 189–221

Contents lists available at ScienceDirect

Artificial Intelligence

www.elsevier.com/locate/artint

Reasoning about discrete and continuous noisy sensors and effectors in dynamical systems

Vaishak Belle [a,b,*], Hector J. Levesque [c]

[a] University of Edinburgh, Edinburgh, United Kingdom
[b] The Alan Turing Institute, London, United Kingdom
[c] University of Toronto, Toronto, Canada

ARTICLE INFO

Article history:
Received 27 October 2016
Received in revised form 26 March 2018
Accepted 5 June 2018
Available online xxxx

Keywords:
Knowledge representation
Reasoning about action
Reasoning about knowledge
Reasoning about uncertainty
Probabilistic logical models
Cognitive robotics

ABSTRACT

Among the many approaches for reasoning about degrees of belief in the presence of noisy sensing and acting, the logical account proposed by Bacchus, Halpern, and Levesque is perhaps the most expressive. While their formalism is quite general, it is restricted to fluents whose values are drawn from discrete finite domains, as opposed to the continuous domains seen in many robotic applications. In this work, we show how this limitation in that approach can be lifted. By dealing seamlessly with both discrete distributions and continuous densities within a rich theory of action, we provide a very general logical specification of how belief should change after acting and sensing in complex noisy domains.

1. Introduction

On numerous occasions it has been suggested that the formalism [the situation calculus] take uncertainty into account by attaching probabilities to its sentences. We agree that the formalism will eventually have to allow statements about the probabilities of events, but attaching probabilities to all statements has the following objections:

1. *It is not clear how to attach probabilities to statements containing quantifiers in a way that corresponds to the amount of conviction people have.*
2. *The information necessary to assign numerical probabilities is not ordinarily available. Therefore, a formalism that required numerical probabilities would be epistemologically inadequate.*

— *McCarthy and Hayes [1].*

Much of high-level AI research is concerned with the behavior of some putative agent, such as an autonomous robot, operating in an environment. Broadly speaking, an intelligent agent interacting with a dynamic and incompletely known world grapples with two special sorts of reasoning problems. First, because the world is *dynamic*, it will need to reason about change: how its actions affect the state of the world. Pushing an object on a table, for example, may cause it to fall on the floor, where it will remain unless picked up. Second, because the world is incompletely known, the agent will need to make do with partial specifications about what is true. As a result, the agent will often need to augment what it believes about the world by performing perceptual actions, using sensors of one form or another.

* Corresponding author at: University of Edinburgh, Edinburgh, United Kingdom.
E-mail addresses: vaishak@ed.ac.uk (V. Belle), hector@cs.toronto.edu (H.J. Levesque).

https://doi.org/10.1016/j.artint.2018.06.003

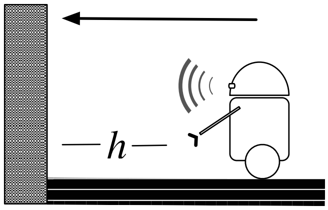

Fig. 1. A simple robot.

For many AI applications, and robotics in particular, these reasoning problems are more involved. Here, it is not enough to deal with incomplete knowledge, where some formula ϕ might be unknown. One must also know which of ϕ or $\neg\phi$ is the more *likely*, and by how much. In addition, both the sensors and the effectors that the agent uses to modify its world are often subject to uncertainty in that they are *noisy.*

To see a very simple example, imagine a robot moving towards a wall as shown in Fig. 1, and a certain distance h from it. Suppose the robot can move towards and away from the wall, and it is equipped with a distance sensor aimed at the wall. Here, the robot may not know the true value of h but may believe that it takes values from some set, say $\{2, \ldots, 11\}$. If the sensor is noisy, a reading of, say, 5 units, does not guarantee that the agent is actually 5 units from the wall, although it should serve to increase the agent's degree of belief in that fact. Analogously, if the robot intends to move by 1 unit and the effector is noisy, it may end up moving by 0.9 units, which the agent does not get to observe. Be that as it may, the robot's degree of belief that it is closer to the wall should increase.

While many proposals have appeared in the literature to address such concerns (cf. penultimate section), very few are embedded in a general theory of action whilst supporting features like disjunction and quantification. For example, graphical models such as Bayesian networks can represent and reason about the probabilistic dependencies between random variables, and how that might change over time. However, it lacks first-order features and a rich account of actions. Relational graphical models, including Markov logic networks [2], borrow devices from first-order logic to allow the succinct modeling of relational dependencies, but ultimately they are purely syntactic extensions to graphical models, and do not attempt to address the deeper issues pertaining to the specification of probabilities in the presence of logical connectives and quantifiers. Building on first-order accounts of probabilistic reasoning [3,4], perhaps the most general formalism for dealing with *degrees of belief* in formulas, and in particular, with how degrees of belief should evolve in the presence of noisy sensing and acting is the account proposed by Bacchus, Halpern, and Levesque [5], henceforth BHL. Among its many properties, the BHL model shows precisely how beliefs can be made less certain by acting with noisy effectors, but made more certain by sensing (even when the sensors themselves are noisy).

The main advantage of a logical account like BHL is that it allows a specification of belief that can be partial or incomplete, in keeping with whatever information is available about the application domain. It does not require specifying a prior distribution over some random variables from which posterior distributions are then calculated, as in Kalman filters, for example [6]. Nor does it require specifying the conditional independences among random variables and how these dependencies change as the result of actions, as in the temporal extensions to Bayesian networks [7]. In the BHL model, some logical constraints are imposed on the initial state of belief. These constraints 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 account using a possible-world model of belief [8–10]. In classical possible-world semantics, a formula ϕ is believed to be true when ϕ holds in all possible worlds that are deemed accessible. In BHL, 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. (Inaccessible worlds are assigned a weight of zero.) To reason about belief change, the BHL model is then embedded in a rich theory of action and sensing provided by the situation calculus [1,11,12]. 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 constraints and these changes in weights.

For example, suppose h is a fluent representing the robot's horizontal distance to the wall in Fig. 1. The fluent h would have different values in different possible worlds. In a BHL specification, each of these worlds might be given an initial weight. For example, a uniform distribution might give an equal weight of .1 to ten possible worlds where $h \in \{2, 3, \ldots, 11\}$. The degree of belief in a formula like $(h < 9)$ is then defined as a sum of the weights, and would lead here to a value of .7. The theory of action would then specify how these weights change as the result of acting (such as moving away or towards the wall) and sensing (such as obtaining a reading from a sonar aimed at the wall). Naturally, the logical language permits weaker specifications, involving disjunctions and quantifiers, and the appropriate behavior would still emerge.

While this model of belief is widely applicable, it does have one serious drawback: it is ultimately based on the addition of weights and is therefore restricted to fluents having discrete finite values. This is in stark contrast to robotics and machine learning applications [13–15], where event and observation variables are characterized by continuous distributions, or perhaps combinations of discrete and continuous ones. There is no way to say in BHL that the initial value of h is any real number drawn from a uniform distribution on the interval $[2, 12]$. One would again expect the belief in $(h < 9)$ to be .7,