

Introducing a Logic for Real-world Agents with Degrees of Belief

Gavin B. Rens and Ivan José Varzinczak

We are developing a new logic (DyMoDeL) for reasoning about dynamic agents with stochastic action, partial observation, *knowledge-of* and *degree-of-belief-in* a state of affairs. Development of the syntax and semantics is almost complete, and we have made progress on how a domain can be formally specified with the logic. Our intention is that the logic presented here forms the foundation of an agent programming language that can generate control policies for decision theoretic agents.

In cognitive robotics, one way to capture the agent’s uncertain knowledge at any time-point, is to have it represent all the states it believes possible as its current belief. A state in this sense, is one way the world could be. A further step to capture uncertain knowledge is to assign a likelihood to each state considered possible so that degrees of belief can be captured.

Now as the robot acts and observes while completing its tasks, it will change its belief accordingly. The robot’s belief reflects its mental state; we say the agent maintains its *belief state* and it *updates* its belief state after every action and observation. Given its current belief state, a robot may want to ask w.r.t. a state of affairs σ , whether it *knows* σ as a matter of fact, whether σ is believed to be merely *possible* or known to be *impossible*, and what the robot’s *degree of belief* in σ is.

Observations and their interaction with actions are dealt with in a new way for logics: When an agent imagines a sequence of actions and observations, and wonders what its attitude towards a state of affairs might be after that sequence, the agent is *projecting*. In such an imagined sequence, the agent may have a clear idea about which actions/observations will be performed/perceived, and in which order. But some actions and observations may be unknown to the agent; it will then want to leave them unspecified during projection. In other words, the robot may also ask itself what its future attitude towards an imagined state of affairs σ will be, after a sequence of action-observation pairs, given some constraints on the number of actions it is willing to perform or given a relaxation of constraints on ‘legal’ actions and percepts.

The logic we introduce in this work is based on two rich formalisms: dynamic logic and the partially observable Markov decision process (POMDP) model.

Several logics can express some of what—we argue—is needed to be expressed; we take from them what we need. However, none of them expresses all that *we* want to express, with the two formalisms we want to use, at the standards we have set for the logic. Moreover, we aim to offer a logical language with which to specify an agent and its interaction with its environment, modeled as a POMDP.

Besides the analyses of decidability and computational complexity, the next task in the development of DyMoDeL is to design a proof theory and determine whether the proof theory is sound and complete. The ultimate goal for DyMoDeL is to extend it into an agent programming language.