REBA: A Refinement-Based Architecture for Knowledge Representation and Reasoning in Robotics

Mohan Sridharan[1] , **Michael Gelfond**[2] , **Shiqi Zhang**[3] , **Jeremy Wyatt**[4]
[1]School of Computer Science, University of Birmingham, UK
[2]Department of Computer Science, Texas Tech University, USA
[3]Department of Computer Science, SUNY Binghamton, USA
[4]School of Computer Science, University of Birmingham, UK
m.sridharan@bham.ac.uk, michael.gelfond@ttu.edu, zhangs@binghamton.edu,
jeremy.l.wyatt@gmail.com

This paper describes REBA, our refinement-based knowledge representation and reasoning architecture for robots that combines the complementary strengths of declarative programming and probabilistic graphical models (Sridharan et al. 2019). It can be viewed as a logician and a statistician communicating through a controller, as shown in Figure 1. REBA is based on tightly-coupled transition diagrams that represent incomplete commonsense knowledge about the domain, and the robot's abilities and beliefs, at two levels of granularity. For any given goal, non-monotonic logical reasoning with the coarse-resolution system description and the system's recorded history results in a sequence of *abstract* actions. Each such abstract action is implemented as a sequence of *concrete* actions by *zooming* to a part of the fine-resolution transition diagram relevant to this abstract action, and probabilistically modeling the non-determinism in action outcomes. REBA makes the following contributions:

1. Action language $\mathcal{AL}_d$ (Gelfond and Inclezan 2013) is extended to support non-Boolean fluents and non-deterministic causal laws, and used to describe the coarse-resolution and fine-resolution transition diagrams of any given domain.
2. The notion of a dynamic domain's history is extended to include *default knowledge* in the initial state, and a model of this history is defined. These definitions are used to define a notion of explanation of unexpected observations, and to provide an algorithm for coarse-resolution planning and diagnostics. This algorithm translates domain knowledge and history into a program of CR-Prolog (Balduccini and Gelfond 2003), an extension to *Answer Set Prolog*, computes answer set of this program, and extracts a plan and an explanation (if needed) from the answer set.
3. A formal definition is provided of one transition diagram being a *weak refinement* of another diagram, and a fine-resolution diagram is defined as a weak refinement of the domain's coarse-resolution transition diagram.
4. A *theory of observations* is introduced and a formal definition is provided of one transition diagram being a *strong refinement* of another. An approach is provided for combining the theory of observations with the weakly-refined fine-resolution transition diagram to obtain a fine-resolution transition diagram that is a strong refinement of the coarse-resolution transition diagram. We establish the tight coupling between the transition diagrams by proving that any coarse-resolution transition can be implemented as a sequence of transitions in the fine-resolution transition diagrams.
5. The *randomization* of the fine-resolution transition diagram is defined. A methods is provided for experimentally collecting statistics and computing the probabilities of the fine-resolution action outcomes and observations.
6. A formal definition is provided for *zooming* to a part of the randomized fine-resolution diagram relevant to any given coarse-resolution (abstract) transition. This definition is used to automate the zoom operation that is invoked to implement each coarse-resolution transition.
7. An algorithm is provided for automatically constructing suitable data structures for the fine-resolution probabilistic implementation of any given abstract action. This algorithm uses probabilistic models of the uncertainty in sensing and actuation, and the zoomed part of the fine-resolution transition diagram. Our implementation uses this algorithm to automatically construct partially observable Markov decision process (Kaelbling, Littman, and Cassandra 1998) models, and uses an approximate solver to compute a policy that is invoked repeatedly to probabilistically execute a sequence of fine-resolution actions that implements the coarse-resolution transition. The corresponding outcomes update the coarse-resolution history for subsequent reasoning.
8. The final and potentially the most important contribution is a general methodology for the design of software components of robots that are re-taskable and robust. This design methodology is based on Dijkstra's view of *step-wise refinement* of a program's specification.

The key advantages of REBA are:

- It substantially simplifies the design process and increases confidence in the correctness of the robot's behavior.
  - Step-wise refinement leads to *separation of concerns* and simplifies testing of the architecture's components.
  - The formal (i.e., mathematical) descriptions of the architecture's components, and of the flow of control and

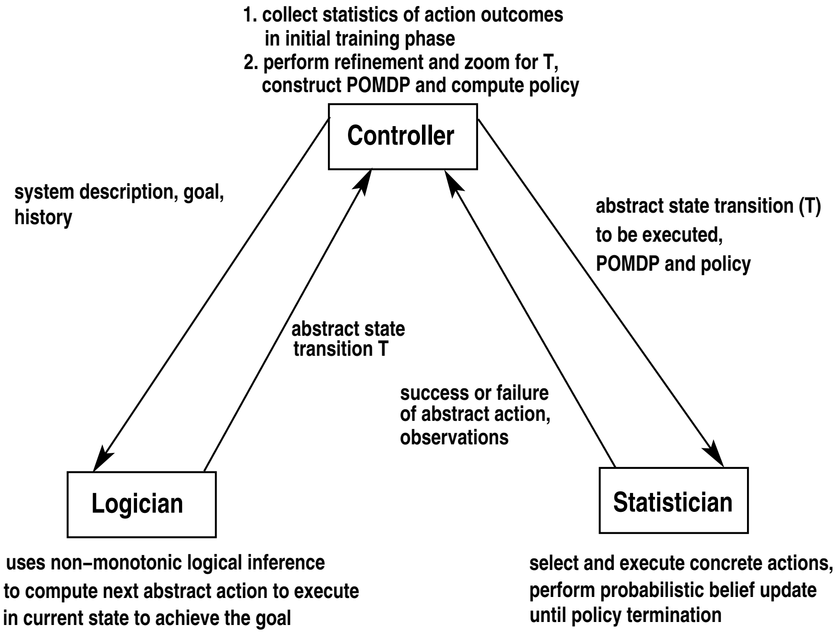

Figure 1: REBA can be viewed as a *logician* and a *statistician* communicating through a *controller*. It combines the complementary strengths of declarative programming and probabilistic reasoning.

information between the components, helps characterize the robot's behavior accurately and *prove correctness* of the algorithms.

  - The domain-independent representations of part of the robot's commonsense knowledge, e.g., theory of observations and strong refinement, enable reuse of these representations on other robots and domains.
  - There is a single framework for inference, planning, diagnostics, and for a quantifiable trade off between accuracy and computational efficiency in the presence of probabilistic models of the uncertainty in sensing and actuation in the physical world.

- It significantly improves the computational efficiency and reliability of planning and action execution on the robot.
  - REBA supports reliable and efficient reasoning with hierarchically-organized knowledge and beliefs at different resolutions. The intentional separation of non-monotonic logical reasoning and probabilistic reasoning is at the heart of the representational elegance, reliability, and inferential efficiency provided by our architecture in complex domains.
  - The *tight coupling* between representation and reasoning at different resolutions, established by formally defining concepts such as refinement and zooming, supports precise reasoning while demonstrating the potential to scale to complex domains by directing attention to the relevant knowledge and observations.
  - Experimental results in simulation and on physical (wheeled) robots deployed in different indoor domains indicate the ability to reason at the sensorimotor level and the cognitive level, in the presence of violation of defaults, noisy observations, and unreliable actions, in more complex domains than was possible before.

Acknowledgments

This work was supported in part by the U.S. Office of Naval Research Science of Autonomy Awards N00014-13-1-0766 (Mohan Sridharan, Shiqi Zhang) and N00014-17-1-2434 (Mohan Sridharan), the Asian Office of Aerospace Research and Development award FA2386-16-1-4071 (Mohan Sridharan), and the EC-funded Strands project FP7-IST-600623 (Jeremy Wyatt). All claims and conclusions reported in the paper are those of the authors.

References

Balduccini, M., and Gelfond, M. 2003. Logic Programs with Consistency-Restoring Rules. In *AAAI Spring Symposium on Logical Formalization of Commonsense Reasoning*.

Gelfond, M., and Inclezan, D. 2013. Some Properties of System Descriptions of AL_d. *Journal of Applied Non-Classical Logics, Special Issue on Equilibrium Logic and Answer Set Programming* 23(1-2):105–120.

Kaelbling, L.; Littman, M.; and Cassandra, A. 1998. Planning and Acting in Partially Observable Stochastic Domains. *Artificial Intelligence* 101:99–134.

Sridharan, M.; Gelfond, M.; Zhang, S.; and Wyatt, J. 2019. REBA: A Refinement-Based Architecture for Knowledge Representation and Reasoning in Robotics. *Journal of Artificial Intelligence Research* 65:87–180. https://jair.org/index.php/jair/article/view/11524.