Extended Abstract: Integrated Commonsense Reasoning and Deep Learning for Transparent Decision Making in Robotics

Mohan Sridharan[1] , **Tiago Mota**[2] , **Ales Leonardis**[1]
[1]Intelligent Robotics Lab, School of Computer Science, University of Birmingham, UK
[2]Electrical, Computer and Software Engineering, The University of Auckland, NZ
m.sridharan@bham.ac.uk, tandrademota@gmail.com, a.leonardis@bham.ac.uk

We summarize work on an architecture that enables a robot to provide on-demand *explanations* of its decisions and beliefs in the form of descriptions comprising relations between relevant domain objects, object attributes, robot attributes, and robot actions. Such "explainability" will help improve the underlying algorithms, establish accountability, and support effective human-robot collaborate. For full details, please see (Mota, Sridharan, and Leonardis 2021).

State of the art robot architectures often include a combination of knowledge-based reasoning methods (e.g., for planning) and data-driven learning methods (e.g., for recognizing objects and events of interest). Providing transparency is particularly challenging in such *integrated robot systems* that require the robot to sense and interact with the physical world, represent and reason with different descriptions of incomplete domain knowledge and uncertainty (e.g., logic-based descriptions of commonsense domain knowledge, probabilistic descriptions of the information from sensors), and incrementally revise its knowledge of domain dynamics (e.g., axioms governing actions and change).

Towards achieving the desired transparency in integrated robot systems, our architecture builds on KR tools and research in cognitive systems that highlights the benefits of coupling different representations, reasoning methods, and learning methods. It combines the complementary strengths of non-monotonic logical reasoning, deep learning, and inductive learning to support the following capabilities:

- Make decisions based on non-monotonic logical reasoning and probabilistic reasoning with incomplete domain knowledge and observations at different resolutions;
- In situations when reasoning is unable to complete the target tasks, automatically identify and use relevant information to learn (deep network) models for these tasks;
- Automatically identify and use relevant information to learn previously unknown axioms encoding constraints and action preconditions and effects;
- Automatically trace the evolution of any given belief or the non-selection of any given action at a given time by inferring the relevant sequence of axioms and beliefs; and
- Exploit the interplay between representation, reasoning, and learning to provide on-demand descriptions of decisions and the evolution of beliefs.

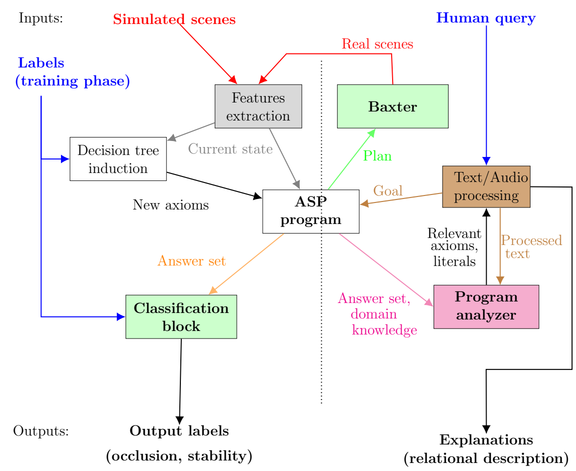

Figure 1: Architecture combines non-monotonic logical reasoning with commonsense domain knowledge, probabilistic reasoning, deep learning, and inductive learning.

These capabilities are evaluated on a simulated and a physical robot manipulating tabletop objects. The robot: (i) computes and executes plans to arrange objects in desired configurations; and (ii) estimates occlusion of objects and stability of object configurations. Figure 1 provides an overview of the architecture, with a Baxter robot performing the tasks. Components to the left of the dashed vertical line combine non-monotonic logical reasoning and deep learning for the chosen tasks. Components to the right of the dashed line expand reasoning and answer questions about decisions and beliefs before, during, or after reasoning and learning.

Our architecture uses Answer Set Prolog (ASP) (Gelfond and Kahl 2014) to encode incomplete domain knowledge that includes object attributes, spatial relations between objects; other domain attributes; features extracted from images of scenes in the domain; and some axioms governing domain dynamics (e.g., prioritized defaults, constraints, causal laws). The robot performs non-monotonic logical reasoning with this knowledge to compute plans to achieve the desired goal and/or to complete the desired estimation tasks (e.g., estimate object occlusion and stability), using probabilistic reasoning when necessary. If ASP-based reasoning is unable to complete the desired tasks, or achieves an incorrect outcome (when ground truth is available), it is

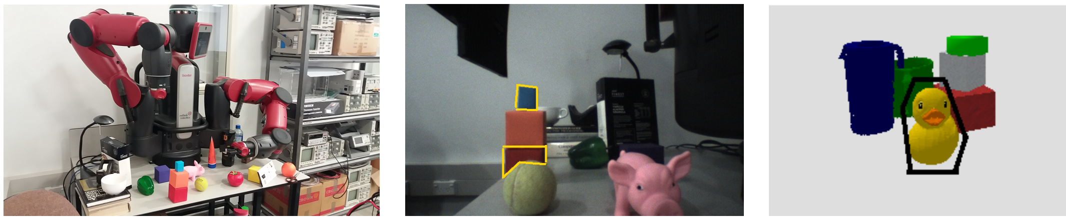

Figure 2: Experimental setup and example scenes: (left) Baxter robot setup; (center) Baxter's camera view; and (right) Simulated scene.

considered to be an indication of missing or incorrect knowledge; this triggers learning. For the estimation tasks, the robot automatically identifies and uses regions of interest (ROIs) in the relevant images to train deep networks that perform these tasks. For both the planning and estimation tasks, information from the relevant ROIs are used with a decision-tree induction algorithm to learn previously unknown axioms and revise existing axioms. The robot also parses the human (verbal or textual) input to identify different types of queries (e.g., descriptive, contrastive, counterfactual). It automatically traces relevant beliefs and axioms to construct relational descriptions that are inserted in templates based on a controlled vocabulary to respond to the queries.

Experimental results indicate the ability to: (i) make decisions reliably and efficiently despite incomplete knowledge and noisy sensor inputs; (ii) incrementally reduce uncertainty in the scene by learning previously unknown constraints, and preconditions and effects of actions; and (iii) construct explanations reliably and efficiently by automatically identifying and reasoning with relevant knowledge.

Execution Example 1. *[Plans, actions, and beliefs]* Consider the image in Figure 2(center) of the tabletop in Figure 2(left), captured by the camera on the robot's arm. The following interaction takes place *after* the robot has executed a plan to move the red cube on the orange cube.

- **Human:** "Please describe the executed plan."
  Baxter: "I picked up the blue cube. I put the blue cube on the table. I picked up the orange cube. I put the orange cube on the table. I picked up the red cube. I put the red cube on the orange cube."
- The human may ask the robot to justify a particular action.
  Human: "Why did you pick up the blue cube at step 0?"
  Baxter: "Because I had to pick up the red cube, and it was below the blue cube."
  This answer is also used to automatically highlight the relevant image regions—see Figure 2(center).
- The human may ask about actions that were not considered.
  Human: "Why did you not put the orange cube on the blue cube?"
  Baxter: "Because the blue cube is small." This answer is based on the single instance (in the executed plan) of putting an orange cube on another cube or surface; it also uses the default knowledge that a large object placed on a small object is typically unstable.
- The robot can be asked to justify particular beliefs.
  Human: "Why did you believe that the red cube was below the blue cube in the initial state?"
  Baxter: "Because I observed the red cube below the blue cube in step zero."

Execution Example 2. *[Learning and explanation]* Even when the robot is unable to achieve the desired object configuration or belief, our architecture enables it to answer questions about its decisions. Consider the following interaction related to the simulated scene in Figure 2(right):

- **Human:** "Please put the pitcher on the duck."
  This action is not executed because the robot learned earlier that placing an object on an irregular surface results in an unstable configuration.

  Human: "Why did you not put the pitcher on the duck?".
  Robot: "Because the duck has an irregular surface."
  The relevant image region is highlighted in Figure 2(right). This example also illustrates how reasoning with learned knowledge helps justify decisions that prevent unfavorable outcomes.

In this scenario, the robot is asked to move the duck on top of the red cube. The computed plan has six steps: pick up the green cylinder, put it on the table, pick up the white cube, put it on the top of the green cylinder, pick up the duck, and put it down on the red cube. Consider the following interaction after the robot has executed this plan:

- **Human:** "Why did you not pick up the green cylinder at step 5?"
  To answer this question, the robot explores the related hypothetical scenario by tracing the evolution of relevant beliefs. It provides the following answer:
  Robot: "Because the white cube was on green cylinder."

  The human may ask for further details:
  Human: "Why did you believe the white cube was on the green cylinder?"
  To answer this question, the robot uses the known causal relationship between the relevant action (*putdown*) and spatial relation (*on*).
  Robot: "Because I put the white cube on the green cylinder at step 4."

References

Gelfond, M., and Kahl, Y. 2014. *Knowledge Representation, Reasoning and the Design of Intelligent Agents*. Cambridge University Press.

Mota, T.; Sridharan, M.; and Leonardis, A. 2021. Integrated Commonsense Reasoning and Deep Learning for Transparent Decision Making in Robotics. *Springer Nature Computer Science* 2(242):1–18.