Knowledge Representation for Autonomous Robotics

Anna Koop ANNA@CS.UALBERTA.CA
Marc G. Bellemare MGBELLEMARE@UALBERTA.CA
Department of Computing Science, University of Alberta, Edmonton, AB T6G 2E8

Abstract

Bridging the gap between the kinds of concepts humans use and the low-level data of robotics is a long-standing goal of Artificial Intelligence research. The Critterbot research program, a new development at the RLAI group, provides us with a domain for pragmatically studying the problem of grounded knowledge representation. With the Critterbot knowledge-representation project, we will be pushing the boundaries of experience-centric knowledge representations towards human-level abstractions. This abstract describes the tools and first steps of the Critterbot research program for knowledge representation.

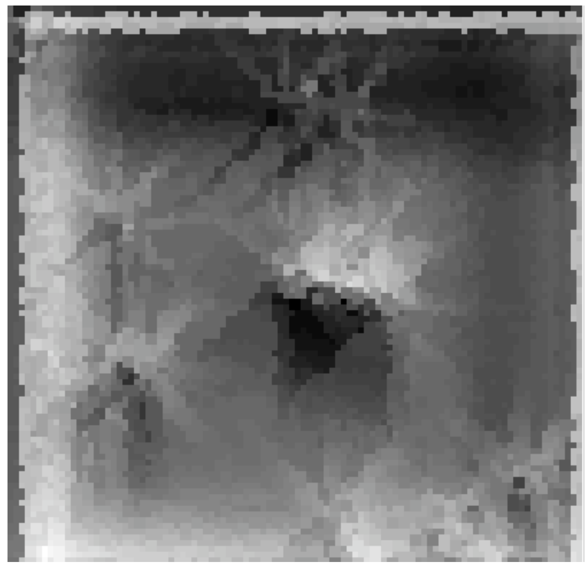

Figure 1. Distance map for points within the agent's representation overlaid on the simulator map.

1. Critterbot Tools and Research Agenda

Our approach to knowledge representation is to ground all concepts in the raw experience of the robot: the input of its sensors and output of its motor controls. This is in contrast to the standard approach to robotics, where systems that correspond to the physics and conceptual framework of the observer must be carefully engineered. We reject this standard framework and attempt instead to uncover patterns that are useful from the agent's perspective. The Critterbot robotic platform has been developed with this aim in mind. The sensory apparatus of the current Critterbot does not include visual data or laser sensors for localization. Instead, the robot has been provided with a wide range of sensors: infrared, light, temperature, sound, acceleration, and other more esoteric sensors. Working in this unusual sensor space will encourage us to focus on the data available to the agent rather than presupposing and encoding an objective location or visual object recognition.

A complementary simulator that allows for flexible environment and sensory specifications is also under development. This will provide a scalable testing ground for subjective knowledge representation. In the simulator environment, agent-centric knowledge can be directly compared to features of the environment. Figure 1 illustrates a subjective mapping comparison between objective position in the simulator and distance within the simulated robot's representation. The shading of each square indicates the distance from the subjective state when the robot was in that location to a reference subjective state. Visualization tools such as this will help us to tease out what can and cannot be encoded in a given subjective representation.

The Critterbot robotic platform is unique in that it is intended for continuous autonomous operation. Although researchers are able to connect to the Critterbot and have it run various programs, the Critterbot will also have an underlying program that is running at all other times, collecting experience and learning about its environment. The idea of lifelong environment discovery relates to, among others, Utgoff and Stracuzzi's work on how knowledge may be accumulated, for example throughout an agent's life (Utgoff & Stracuzzi, 2002). The Critterbot will be responsible for charging its batteries and will be able to engage in exploratory behavior, living an independent "life" while not under explicit direction. Having a robotic agent with a long life and ample opportunity to learn about the daily and weekly cycles in the lab is an novel benefit of the Critterbot project.

The Critterbot project also contains a flexible framework for defining reward functions, while the action and sensation space is consistent for the robot or a specific simulated agent. Because of this consistency, we believe that generalization will be possible and some kinds of subjective knowledge will be useful across a wide range of tasks. Konidaris and Barto's *agent space* is one way of describing this kind of subjective knowledge (Konidaris & Barto, 2007). The Critterbot framework provides a clear method for empirically evaluating the performance of candidate representations on a range of tasks.

2. First Steps for Autonomous Knowledge

We have previously proposed that abstract knowledge can be defined in experience terms through the frameworks of predictions and options (Koop, 2007). The Critterbot research program will allow for the first full-scale implementation of this joint framework for knowledge representation. This implementation will be initially developed through three stages.

In the first stage of the project, we will examine methods for learning the parameters of option-conditioned predictions. Options and prediction targets will be hand-coded, and the parameters of a linear or semi-linear function learned online by the agent. Various learning methods from the reinforcement learning and supervised learning communities will be compared. Two initial candidates are the gradient temporal-difference learning algorithm and the optimal reverse prediction method (Sutton et al., 2009; Xu et al., 2009).

Once the predictions have been learned, the next step will be to evaluate the hypothesis that such predictions are "useful" knowledge. The "usefulness" of knowledge can be evaluated in two ways. The first evaluation is through studying the performance benefit of using such knowledge with respect to a set of reinforcement learning tasks. The second way in which we will evaluate our knowledge representation is by studying the change in error when prediction questions from within a fixed set. We will compare the learning power of agents that have predictive knowledge in the agent state to the learning power of agents with simpler state representations in a variety of frameworks, and this will give us a test of the hypothesis that subjective knowledge is useful.

3. Conclusion

Grounding knowledge in data readily accessible to an intelligent agent is one of the most important challenges facing autonomous robotics. With the Critterbot platform, we have new tools for teasing apart the issues of subjective representation. We will be testing the use of option-conditioned predictions for knowledge representation in a complex and extendible domain, and through this advance AI towards human-level intelligence.

References

Konidaris, G., & Barto, A. (2007). Building portable options: Skill transfer in reinforcement learning. *Proceedings of the 20th International Joint Conference on Artificial Intelligence* (pp. 895–900).

Koop, A. (2007). Understanding experience: Temporal coherence and empirical knowledge representation. Master's thesis, University of Alberta.

Sutton, R. S., Szepesvari, C., & Maei, H. R. (2009). A convergent O(n) algorithm for off-policy temperal-difference learning with linear function approximation. *Advances in Neural Information Processing Systems 21.*

Utgoff, P. E., & Stracuzzi, D. J. (2002). Many-layered learning. *Neural Computation*, *14*, 2497–2529.

Xu, L., White, M., & Schuurmans, D. (2009). Optimal reverse prediction: A unified perspective on supervised, unsupervised and semi-supervised learning. *To appear in Proceedings of the 20th International Joint Conference on Artificial Intelligence.*