Transfer of Learned Knowledge in Lifelong Learning Agents

Joseph O'Sullivan

Problem: The cost of gathering and labelling data is remaining constant, while computing power is becoming cheaper. Recognizing this, the notion of a lifelong learning agent has been proposed. A lifelong learning agent is a system that learns knowledge and, when thereafter learning novel tasks, can use this knowledge to increase the accuracy of what is learned or to reduce the number of examples necessary for learning. Many open problems exist: can an agent exploit multiple sources of learned knowledge, to what degree will an agent benefit from different types of learned knowledge, what types of knowledge should be made available, how should the agent adapt as a new task arrives, how might the order of task arrival impact learning, and of course, how should such an agent be built?

Impact: A lifelong learning agent offers the promise of learning from substantially fewer examples. Currently, we have shown situations were the number of examples required to learn a task have been halved. Understanding how and where a lifelong agent can be applied would allow systems to be built that efficiently learn novel tasks.

State of the Art: This problem area is still quite young and fertile, and so there is quite a wide diversity of approaches in the literature. In the majority of previous work on agents, as the agent ages and further new tasks arise, the agent does not improve its ability to learn those tasks. Of the wide variety of underlying mechanisms proposed for creating a lifelong long agent, few have been applied to a complete system, and not enough is understood about how the mechanisms scale as an agent ages. There have been several different types of lifelong learning agents systems proposed – including variants of reinforcement learning and of high level planning systems. Generally, these systems either make abstracting assumptions about the agent input, or are limited in scope.

Approach: I propose that an agent can be constructed which learns knowledge and exploits that knowledge to effectively improve further learning by reducing the number of examples required to learn. I am studying the transfer of learned knowledge by supervised life-long learning agents within a neural network based architecture capable of increasing capacity with the number of tasks faced. So far, preliminary work has been carried out in controlled settings, and an appropriate architecture has been outlined. This work has confirmed that learned knowledge can reduce the number of examples required to learn novel tasks and that combining previously separate mechanisms can yield a synergistic improvement on learning ability. It has shown that capacity can be expanded as new tasks arise over time. The underlying mechanisms that are being used for transfer include learning useful representations for a specific domain, and learning domain specific models for use in further learning. Our initial work has explored the benefits of these mechanisms separately, and how to extend these methods to operate concurrently, and as the agent ages.

The hypothesizes are being studied experimentally both in simulated domains, and will be verified on the robot test-bed in the Learning Robot Lab; Xavier which we designed and constructed internally within the laboratory in 1993, and Amelia, a commercial descendent of Xavier.

Future Work: We are currently building upon the underlying mechanisms to generate a tool-box for building curriculums. That is, given that a lifelong learning agent exists, what is the best things to teach the agent, and the best way to teach them, so that we expect to learn novel tasks efficiently. We speculate that the order in which tasks arise can be exploited with a graded curriculum.

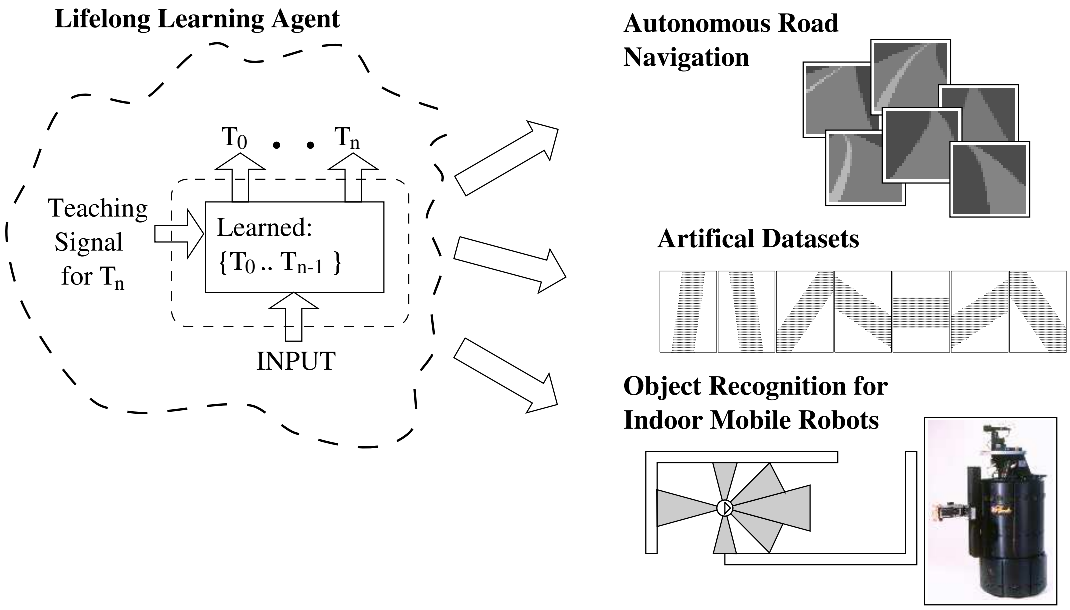

Figure 1: The lifelong learning agent architecture under investigation is being evaluated in a variety of simulated and real world domains. When learning a novel task T_n, a lifelong learning agent can exploit knowledge learned from previous tasks $T_0..T_{n-1}$.

Selected References:

[1] Rich Caruana and Joseph O'Sullivan. Backprop nets do unsupervised clustering of outputs. Submitted to ICANN98, 1998.

[2] Rich Caruana and Joseph O'Sullivan. Multitask pattern recognition for autonomous robots. Submitted to IROS98, 1998.

[3] Joseph O'Sullivan. The CMU Learning Robot Laboratory Data Toolkit. In *Proceedings of MLC-Colt Workshop on Robot Learning*, Rutgers University, New Brunswick, N.J, July 1994.

[4] Joseph O'Sullivan. Transfer of learned knowledge in life-long learning agents. PhD Thesis Proposal, Carnegie Mellon University, Feb 1997.

[5] Joseph O'Sullivan. Sequential multitask learning for a lifelong learning agent. Submitted to ICML98, 1998.

[6] Joseph O'Sullivan, Karen Haigh, and G. D. Armstrong. Xavier - the Manual v0.3. Revised April 1997, Carnegie Mellon University, School of Computer Science, `http://www.cs.cmu.edu/~Xavier`, March 1994.

[7] Joseph O'Sullivan, Tom M. Mitchell, and Sebastian B. Thrun. Explanation Based Learning for Mobile Robot Perception. In Katsushi Ikeuchi and Manuela Veloso, editors, *Symbolic Visual Learning*. Oxford University Press, 1996.

[8] Reid Simmons, Richard Goodwin, Karen Haigh, Sven Koenig, and Joseph O'Sullivan. A modular architecture for office delivery robots. In *The First International Conference on Autonomous Agents*, Feb 1997.

[9] Reid G. Simmons, Richard Goodwin, Karen Zita Haigh, Sven Koenig, Joseph O'Sullivan, and Manuela M. Veloso. Xavier: Experience with a layered robot architecture. *Intelligence*, 1998. A shorter version appears as [8].

[10] Sebastian Thrun and Joseph O'Sullivan. Clustering learning tasks and the selective cross-task transfer of knowledge. In S. Thrun and L.Y. Pratt (eds), editors, *Learning To Learn*, chapter 10. Kluwer Academic Publisher, 1998. A shorter version appears as [11].

[11] Sebastian B. Thrun and Joseph O'Sullivan. Discovering Structure in Multiple Learning Tasks: The TC Algorithm. In *Proceedings of ICML*, Torino, Italy, July 1996.