Machine Lifelong Learning: Challenges and Benefits for Artificial General Intelligence

Daniel L. Silver

Jodrey School of Computer Science
Acadia University
Wolfville, Nova Scotia, Canada B4P 2R6
danny.silver@acadiau.ca

Abstract. We propose that it is appropriate to more seriously consider the nature of systems that are capable of learning over a lifetime. There are three reasons for taking this position. First, there exists a body of related work for this research under names such as constructive induction, continual learning, sequential task learning and most recently learning with deep architectures. Second, the computational and data storage power of modern computers are capable of implementing and testing *machine lifelong learning* systems. Third, there are significant challenges and benefits to pursuing programs of research in the area to AGI and brain sciences. This paper discusses each of the above in the context of a general framework for machine lifelong learning.

1 Introduction

Over the last 25 years there have been significant advances in machine learning theory and new machine learning algorithms based on that theory. However, there has been comparatively little work on systems that are able to learn a variety of tasks over an extended period of time. We propose that it is now appropriate to more seriously consider the nature of systems that are capable of learning over a life time. In accord with [13], we call these *machine lifelong learning* systems.

There are three reasons for feeling the time is right to more vigorously explore lifelong learning systems. First, there exists a body of related work that provides a starting point for research under names such as constructive induction, incremental and continual learning, sequential task learning, and most recently learning with deep architectures. Second, the computational and data storage power of modern computers are capable of implementing and testing lifelong learning systems. Third, there are significant challenges and benefits to pursuing programs of research in the area to AGI and brain sciences. This paper presents a general framework for machine lifelong learning and then discusses each of the above reasons for further research.

J. Schmidhuber, K.R. Thórisson, and M. Looks (Eds.): AGI 2011, LNAI 6830, pp. 370–375, 2011.

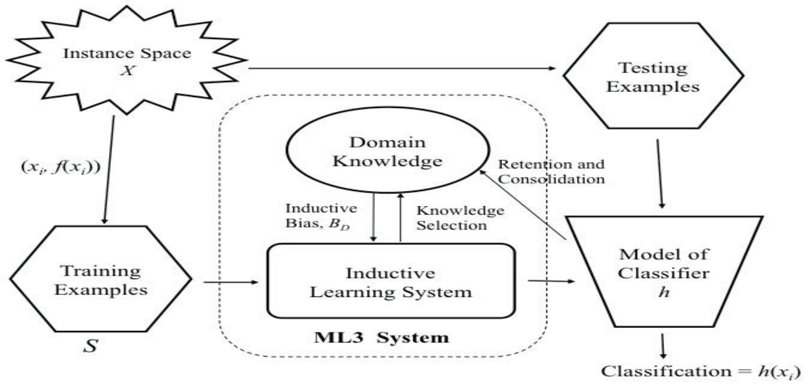

Fig. 1. A framework for machine lifelong learning

2 A Framework for Machine Lifelong Learning

The constraint on a learning system's hypothesis space, beyond the criterion of consistency with the training examples, is called *inductive bias* [3]. Inductive bias is essential for the development of a hypothesis with good generalization from a practical number of examples. Ideally, a lifelong learning system can select its inductive bias to tailor the preference for hypotheses according to the task being learned [15].

Figure 1 provides a general framework for a machine lifelong learning (ML3) approach that uses knowledge of the task domain as a source of inductive bias [8]. As with a standard inductive learner, training examples (supervised and possibly unsupervised) are used to develop a hypothesis of a classification task. However, unlike a standard learning system, knowledge from each hypothesis is saved in a long-term memory structure called *domain knowledge*. When learning a new task, aspects of domain knowledge are selected to provide a beneficial inductive bias to the learning system. The result is a more accurate hypothesis developed in a shorter period of time. The method relies on the transfer of knowledge from one or more prior secondary tasks, stored in domain knowledge, to the hypothesis for a new primary task. The problem of selecting an appropriate bias becomes one of selecting the most related knowledge for transfer. A machine lifelong learning system typically has short-term transfer and long-term retention learning phases. Although two phases of learning may not be necessary, it is frequently required so as to properly consolidate the hypothesis of a new task into long-term domain knowledge.

3 Related Work

Several prior research efforts have considered systems that learn domains of tasks over extended periods of time. In particular, progress has been made in machine learning that exhibit aspects of knowledge retention and inductive transfer. These represent advances in inductive modeling that move beyond *tabula rasa* learning and toward machines capable of lifelong learning [13].

Utgoff and Mitchell wrote in 1983 about the importance of *inductive bias* to concept learning from practical sets of training examples [14]. They theorized that learning systems should conduct their own search for an appropriate inductive bias using knowledge such as that of related tasks. They proposed a system that could shift its bias by adjusting the operations of the modeling language.

In the mid 1980s Michalski introduced the theory of *constructive inductive learning* to cope with learning problems in which the original representation space is inadequate for the problem at hand [2]. New knowledge is hypothesized through two interrelated searches: (1) a search for the best representational space for hypotheses and (2) a search for the best hypothesis within the current representational space. The underlying principle is that new knowledge is easier to induce if search is done using the *right* representation.

In 1989 Solomonof began work on *incremental learning* [11]. His system was primed on a small, incomplete set of primitive concepts, that are able to express the solutions to the first set of simple problems. When the machine learns to use these concepts effectively it is given more difficult problems and, if necessary, additional primitive concepts needed to solve them, and so on.

In the mid 1990s, Thrun and Mitchell worked on a lifelong learning approached they called *explanation-based neural networks* [12]. EBNN is able to transfers knowledge across multiple learning tasks. When faced with a new learning task, EBNN exploits domain knowledge of previous learning tasks (back-propagation gradients of prior learned tasks) to guide the generalization of the new one. As a result, EBNN generalizes more accurately from less data than comparable methods.

Since 1995, Silver *et al.* have proposed several variants of *sequential learning and consolidation systems* using standard back-propagation neural networks [9,10]. A system of two multiple task learning networks is used; one for short-term learning using task rehearsal to selectively transfer prior knowledge, and a second for long-term consolidation using task rehearsal to overcome the stability-plasticity problem. *Task rehearsal* is an essential part of this system. After a task has been successfully learned, its hypothesis representation can saved. The saved hypothesis can be used to generate virtual training examples so as to rehearse the prior task in parallel when learning a new task. It is through the rehearsal of previously learned tasks within the shared representation of the neural network that knowledge is transferred to the new task. Similarly, [9] the knowledge of a new task can be consolidated into a large domain knowledge network without loss of existing task knowledge by using task rehearsal to maintain the function accuracy of the prior tasks while the representation is modified to accommodate the new task.

In 1997, Ring proposed a lifelong learning approach called *continual learning* that builds more complicated skills on top of those already developed both incrementally and hierarchically [4]. He presents a system that can efficiently solve reinforcement-learning tasks and can then transfer its skills to related but more complicated tasks.

Rivest and Schultz proposed *knowledge-based cascade-correlation* neural networks in the late 1990s [5]. The method extends the original cascade-correlation approach, by selecting previously learned sub-networks as well as simple hidden units. In this way it is able to use past learning to bias new learning.

Recent research into the learning of *deep architectures* of neural networks can be connected to lifelong learning [1]. Layered neural networks of unsupervised Restricted Boltzman Machine and auto-encoders have been shown to efficiently develop hierarchies of features that capture regularities in their respective inputs. When used to learn a variety of class categories, these networks develop layers of common features similar to that seen in the visual cortex of humans.

4 Current Computational and Data Storage Capacity

The number of transistors that can be placed cheaply on an integrated circuit has doubled approximately every two years since 1970. This trend is expected to continue until the foreseeable future, with some expecting the power of computing systems to move to a log scale as computing systems increasingly use multiple processing cores. We are now at a point where a lifelong learning system focused on a constrained domain of tasks (*e.g.* medical diagnosis, product recommendation) is computationally tractable in terms of both computer memory and processing time.

As an example, massively parallel data processing engines now exist that are capable of competing with humans in real-time question-answer problems. This was recently witnessed on the Jeopardy television game show in February, of 2011. Watson consisted of 90 IBM server computers, each with four 8-core processors. It used 15 terabytes (220 million text pages) of rapid access memory and divided its tasks into thousands of stand-alone jobs distributed among 80 teraflops (1 trillion operations/second) of parallel processing power. Given that much of machine learning is search, platforms such as the one used by Watson are well suited to the challenges of lifelong learning systems. It would be important to note that Watson's success was in part due to advances in machine learning methods.

5 Challenges and Benefits

There are a number of challenges for and potential benefits from new research programs in machine lifelong learning. The following captures several of these.

There is strong evidence that transfer learning from prior related knowledge is beneficial when learning a new task [5,10,12]. Experimental results indicate that effective learning excels under functional transfer whereas efficient learning requires representation transfer [7]. Recent work has also shown the benefit of unsupervised training using many unlabelled examples as a source of inductive bias for supervised learning [1].

Machine lifelong learning provides an opportunity to acquire and take advantage of related knowledge. However, there are many challenging problems; for example, a lifelong learning system must weigh the relevance and accuracy of retained knowledge along side that of the available training examples for a new task. Theories on how to select inductive bias and modify the representational space of hypotheses [11] will be of significant value to AGI and brain science.

Mechanisms that can effectively and efficiently retain learned knowledge over time will suggest new approaches to common knowledge representation. In particular, methods of overcoming the stability-plasticity problem so as to integrate new knowledge into existing knowledge are of value to researchers in AI, cognitive science and neuroscience [9]. Efficient long-term retention of learned knowledge should cause no loss of prior task knowledge, no loss of new task knowledge, and an increase in the accuracy of old tasks if the new task being retained is related. Furthermore, the knowledge representation approach should allow a lifelong learner to efficiently select the most effective prior knowledge for inductive transfer during short-term learning.

A lifelong learning system should facilitate the practice of a task such that the generalization accuracy of the long-term hypothesis for the task increases. But how can a lifelong learning system determine from the training examples that it is practicing a task it has previously learned versus learning a new but closely related task. Related work suggests that a system should not be explicit in this determination [6,10]; rather, the similarity of a set of training examples to that of prior domain knowledge should be implicit; each training example should be able to draw upon those aspects of domain knowledge that are most related. This suggests that domain knowledge should be seen as continuum as apposed to a set of disjoint tasks. A theory of how best to practice tasks will be useful to the fields of AI, psychology and education.

Scalability is often the most difficult and important challenge for computer scientists. A machine lifelong learning system must be capable of scaling up to large numbers of inputs, outputs, training examples and learning tasks. Preferably, the space and time complexity of the learning system grows polynomially in all of these factors.

Software agents and robots will make good use of lifelong learning systems, or at least provide useful test platforms for empirical studies [12]. Agents and robots will naturally encounter new examples of problems periodically, providing opportunities to test the practice and consolidation of task knowledge.

The study of lifelong learning systems will provided insight into curriculum and training sequences that are beneficial for both humans and machines [11,4]. This will be beneficial to robot and software agent training and will likely lead to the confirmation of and advances in human educational curriculum.

Finally, research into machines that can learn over a lifetime involves laborious repeated studies of lengthy sequences of problems. This is tough but rewarding work that will become less labor intensive as experimental methods develop.

References

1. Bengio, Y.: Learning deep architectures for ai. Foundations and Trends in Machine Learning 2(1), 1–127 (2009)
2. Michalski, R.S.: Learning = inferencing + memorizing. In: Foundations of Knowledge Acquistion: Machine Learning, pp. 1–41 (1993)
3. Mitchell, T.M.: The need for biases in learning generalizations. In: Shavlik, J.W., Dietterich, T.G. (eds.) Readings in Machine Learning, pp. 184–191 (1980)
4. Ring, M.B.: Child: A first step towards continual learning. In: Machine Learning, pp. 77–104 (1997)
5. Shultz, T.R., Rivest, F.: Knowledge-based cascade-correlation: using knowledge to speed learning. Connect. Sci. 13(1), 43–72 (2001)
6. Silver, D.L., Alisch, R.: A measure of relatedness for selecting consolidated task knowledge. In: Proceedings of the 18th Florida Artificial Intelligence Research Society Conference (FLAIRS 2005), pp. 399–404 (2005)
7. Silver, D.L., Mercer, R.E.: The parallel transfer of task knowledge using dynamic learning rates based on a measure of relatedness. Connection Science Special Issue: Transfer in Inductive Systems 8(2), 277–294 (1996)
8. Silver, D.L., Mercer, R.E.: The task rehearsal method of life-long learning: Overcoming impoverished data. In: Advances in Artificial Intelligence, 15th Conference of the Canadian Society for Computational Studies of Intelligence (AI 2002), pp. 90–101 (2002)
9. Silver, D.L., Poirier, R.: Sequential consolidation of learned task knowledge. In: Canadian AI 2004. Lecture Notes in AI, pp. 217–232 (2004)
10. Silver, D.L., Poirier, R., Currie, D.: Inductive tranfser with context-sensitive neural networks. Machine Learning 73(3), 313–336 (2008)
11. Solomonoff, R.J.: A system for incremental learning based on algorithmic probability. In: Probability, Proceedings of the Sixth Israeli Conference on Artificial Intelligence, Computer Vision and Pattern Recognition, pp. 515–527 (1989)
12. Thrun, S.: Explanation-Based Neural Network Learning: A Lifelong Learning Approach. Kluwer Academic Publishers, Boston (1996)
13. Thrun, S.: Lifelong learning algorithms. In: Learning to Learn, pp. 181–209. Kluwer Academic Publishers, Dordrecht (1997)
14. Utgoff, P.E.: Adjusting bias in concept learning. In: Proceedings of IJCAI-1983, pp. 447–449 (1983)
15. Utgoff, P.E.: Machine Learning of Inductive Bias. Kluwer Academc Publisher, Boston (1986)