# On Common Ground: Neural-Symbolic Integration and Lifelong Machine Learning

**Daniel L. Silver**
Acadia University
Nova Scotia, Canada
danny.silver@acadiau.ca

## Abstract

Research efforts in Neural-Symbolic Integration and Lifelong Machine Learning have taken place with limited interactions over the last 20 years. These two areas of artificial intelligence share common ground and yet have much to learn from each other. This paper provides background, particularly on Lifelong Machine Learning, and presents a number of common areas of investigation with Neural-Symbolic Integration. It then invites researchers in the two fields to better understand some of the differences in motives and objectives for the purpose of jointly advancing machine learning, knowledge representation and reasoning.

## 1 Introduction

Recent work on unsupervised learning using deep learning architectures has given rise to new ideas on how knowledge of the world is learned, consolidated, and then used for future learning and reasoning [Bengio, 2009; Le *et al.*, 2012]. This is bringing together research in the areas of machine learning and knowledge representation that have traditionally been pursued separately. Specifically, research efforts in Neural-Symbolic Integration [Garcez *et al.*, 2009; Bader and Hitzler, 2005] and Lifelong Machine Learning [Silver *et al.*, 2013] have taken place with only a few points of interaction over the last 20 years. This paper presents several areas of common ground for these areas where joint work has taken place and further collaboration is possible. The paper also points out fundamental differences in the motives and objectives of NSI versus LML that need to be understood by researchers as we move forward. In particular, we propose that joint research has the potential to make serious advances on a significant problem in artificial intelligence - the learning of *common background knowledge* that can be used for future learning and reasoning.

## 2 Background

### 2.1 Neural-Symbolic Integration

Neural-symbolic integration, or NSI, considers hybrid systems that integrate neural networks and symbolic logic. The goal of NSI is to take advantage of the best of symbolic and connectionist paradigms of artificial intelligence for both learning and reasoning [Garcez *et al.*, 2002; Garcez and Gabbay, 2004].

NSI seeks to make use of the learning capacities of neural network models and the reasoning capacities of logic [Garcez *et al.*, 2014]. In a NSI system, neural networks provide the machinery for parallel computation and robust learning, while symbolic logic provides knowledge representation and reasoning. The retention of symbolic knowledge of learned models can be used for transfer learning, explanation of the models to humans, and use of the knowledge by other systems. Neural-symbolic systems have application in knowledge acquisition where a system learns a complex model and then needs to reason about what has been learned in order to respond to a new situation.

NSI has at least three major areas of investigation. First there is the use of connectionist systems for symbolic knowledge representation, reasoning and learning. These systems include comparisons with purely-symbolic and purely-connectionist models, and the representation of relational, first-order and modal logics for higher-order reasoning [Garcez *et al.*, 2014]. A second area of research is the extraction of high-level concepts and knowledge from complex networks. The focus here is on efficient and effective knowledge extraction from very large networks for the purpose model comprehension, validation, maintenance and transfer learning. The third area of investigation is the design and development of applications in areas such as vision, robotics, intelligent agents, and simulation.

### 2.2 Lifelong Machine Learning

Lifelong Machine Learning, or LML, is concerned with the persistent and cumulative nature of learning [Thrun, 1996b]. LML considers systems that can learn many tasks over a lifetime from one or more domains. An LML system must efficiently and effectively retain the knowledge it has learned and transfer that knowledge to more efficiently and effectively learn new tasks through the transfer of knowledge [Silver *et al.*, 2013]. The following sections provide an overview of prior LML research in all areas of machine learning - supervised, unsupervised and reinforcement learning.

#### Supervised Learning

As early as the mid 1980s Michalski and Solomonoff had theories on *constructive inductive learning* [Michalski, 1993] and *incremental learning* [Solomonoff, 1989]. In the mid 1990s, Thrun and Mitchell worked on *explanation-based neural networks* [Thrun, 1996a] and applied EBNN transfer learning to autonomous robot learning when a multitude of control learning tasks are encountered over an extended period of time [Thrun and O'Sullivan, 1995].

Since 1995, Silver *et al.* have proposed variants of *sequential learning and consolidation systems* using standard back-propagation neural networks [Silver and Poirier, 2004; Silver *et al.*, 2008]. A method called *task rehearsal* is an essential part of these systems. After a task has been successfully learned, its hypothesis representation is saved. The saved hypothesis can be used to generate virtual training examples so as to rehearse the prior task when learning a new task. Knowledge is transferred to the new task through the rehearsal of previously learned tasks within the shared representation of the neural network. Similarly, 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 functional accuracy of the prior tasks while the representation is modified to accommodate the new task.

In the late 1990s, Rivest and Schultz proposed *knowledge-based cascade-correlation* neural networks [Shultz and Rivest, 2001]. The method extends the original cascade-correlation approach, by selecting previously learned sub-networks as well as simple hidden units. In this way the system is able to use past learning to bias new learning.

### Unsupervised Learning

Transfer in unsupervsied learning is almost as old as that of supervised learning. In the mid 1980s, Carpenter and Grossberg proposed ART (Adaptive Resonance Theory) neural networks to overcome the stability-plasticity problem of forgetting previous learned data concepts [Grossberg, 1987].

Raina *et al.* proposed the *Self-taught Learning* method to build high-level features using unlabeled data for a set of tasks [Raina *et al.*, 2007]. The authors used the features to form a succinct input representation for problems such as image and webpage classification.

Recent research into *deep learning architectures* of neural networks can be connected to LML [Bengio, 2009]. Layered neural networks of unsupervised Restricted Boltzman Machine auto-encoders have been shown to efficiently develop hierarchies of features that capture statistical 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. Le *et al.* have used deep learning methods to build high-level features for large-scale applications by scaling up the dataset, the model and the computational resources [Le *et al.*, 2012]. By using millions of high resolution images and very large neural networks, their system effectively discover high-level concepts like the presence of a cat's face in an image. Experimental results show that their network can use its learned features to achieve a significant improvement in image classification performance over state-of-the-art methods.

### Reinforcement Learning

Several reinforcement learning researchers have considered LML systems. 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 [Ring, 1997].

Tanaka and Yamamura proposed a lifelong reinforcement learning method for autonomous-robots by treating multiple environments as multiple-tasks [Tanaka and Yamamura, 1999].

Sutton *et al.* suggest that learning should continue during an agent's operations since the environment may change making prior learning insufficient [Sutton *et al.*, 2007]. An agent is proposed that adapts to different local environments when encountering different parts of its world over an extended period of time.

### Moving Beyond Learning Algorithms

Many machine learning researchers are calling for a move beyond the development of inductive learning algorithms and onto the design of systems that learn, retain and use knowledge over a lifetime. In [Silver *et al.*, 2013] we cite the following reasons for a call for wider research into LML systems.

*Selective Inductive Bias is Essential to Learning.* The constraint on a learning system's hypothesis space, beyond the criterion of consistency with the training examples, is called *inductive bias* [Mitchell, 1980]. Utgoff wrote in 1983 about the importance of inductive bias to concept learning from practical sets of training examples and the need for learning systems to select bias [Utgoff, 1983]. The AI community has come to accept the futility of searching for a universal machine learning algorithm [Wolpert, 1996]. LML systems that retain and selectively use prior knowledge as a source of inductive bias promotes this perspective.

*Theoretical Advances in ML and KR.* Thrun proposed "The acquisition, representation and transfer of domain knowledge are the key scientific concerns that arise in lifelong learning" [Thrun, 1997]. Knowledge representation will play an important role in the development of LML systems. More specifically, the interaction between knowledge retention and knowledge transfer will be key to the design of intelligent agents that learn many things over an extended period.

*Practical Agents/Robots Require LML.* Advances in autonomous robotics and intelligent agents that run on the web or in mobile devices present opportunities for employing LML systems. Robots such as those that go into space or travel under the sea must learn to recognize objects and make decisions over extended periods of time and varied environmental circumstances. The ability to retain and use learned knowledge is very attractive to the researchers designing these systems. Similarly, software agents on the web or in our mobile phones would benefit from the ability to learn more quickly and more accurately as they are challenged to learn new but related tasks from small numbers of examples.

*Increasing Capacity of Computers.* Advances in modern computers provide the computational power for implementing and testing LML systems. The number of transistors that can be placed cheaply on an integrated circuit has doubled

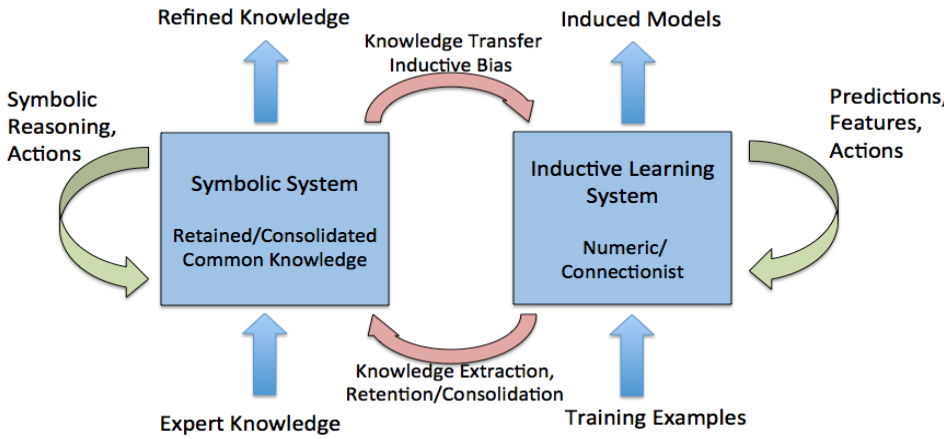


Figure 1: An integrated framework for Neural-Symbolic Integration and Lifelong Machine Learning.

approximately every two years since 1970. This trend is expected to continue into the foreseeable future, as computing systems increasingly use multiple processing cores. We are now at a point where an LML system focused on a constrained domain of tasks (*e.g.* product recommendation) is computationally tractable in terms of both computer memory and processing time [Le *et al.*, 2012].

## 3 Common Ground

Figure 1 is based on a diagram found in [Bader and Hitzler, 2005]. The modifications show how the NSI framework can integrate nicely with the LML framework presented in [Silver *et al.*, 2013]. The integrated framework is meant to encompass the various methods of machine learning (supervised, unsupervised, or reinforcement) and the various symbolic systems. Expert knowledge can be provided by the user to the symbolic system where it is retained and/or consolidated with existing knowledge. This store of common background knowledge can be used by the inductive learning system as a source of knowledge transfer, or inductive bias. This bias can come in representational form, such as an initial set of neural network weights from which to start learning, or in functional form, such as a set of examples for a secondary task for multiple task learning [Shavlik, 1992; Silver and Mercer, 2002]. The inductive learning system develops a hypothesis, or model, using the transferred knowledge and training examples provided by the user. The result is a more accurate model developed in a shorter period of time. The knowledge learned in these models can be extracted back to the symbolic system and retained (or consolidated) for symbolic reasoning, used by another system, or for explanation to the user.

The following sections discuss areas of common ground shared by NSI and LML research, of which there has already been some joint work.

### 3.1 Choice of Machine Learning to Use

An open question for both NSI and LML is which approach to machine learning or combination of approaches works best in the context of knowledge extraction and symbolic reasoning. Supervised learning continues to dominate NSI research [Bader and Hitzler, 2005], with some work being done using reinforcement and unsupervised learning including the extraction of symbolic logic from deep believe networks [Tran and Garcez, 2012]. Supervised learning has traditionally dominated LML, but over the last five years unsupervised learning has received a considerable amount of attention. For example, recent work has shown the benefit of unsupervised training using many unlabelled examples to generate new encodings (features) of the examples for use in supervised learning [Bengio, 2009]. Others feel that reinforcement learning is the only true method of developing predictive models [Sutton *et al.*, 2007]. The choice of machine learning algorithm and representation will dramatically affect the structure and function of NSI and LML systems. Combinations of approaches may be helpful for learning individual tasks but will prove challenging for knowledge consolidation and aspects of NSI.

### 3.2 Training Examples versus Prior Knowledge

Both NSI and LML systems must weigh the relevance and accuracy of retained knowledge along-side the information resident in the available training examples for a new task. An estimate of the sample complexity of the new task will play a role here. The relative accuracy level of prior knowledge must also be considered. Theories of how to selectively transfer common knowledge in combination with existing training examples is of significant value to NSI and LML research.

### 3.3 Effective and Efficient Knowledge Retention

Mechanisms that can effectively and efficiently retain knowledge over time will suggest new approaches to common knowledge representation. In particular, methods of integrating new knowledge into existing knowledge are of value to researchers in NSI and LML [Shavlik, 1992; Silver and Poirier, 2004]. Efficient long-term retention of symbolic or learned knowledge should cause no loss of prior task knowledge and increase the accuracy of such knowledge if the new task being retained is related. Furthermore, the knowledge representation approach should allow the NSI or LML system to efficiently select the most effective prior knowledge for transfer during new learning.

In general, research in NSI and LML will see theories of transfer learning and knowledge representation influence and affect each other. A combined research agenda has the potential to make serious advances on a significant AI problem - the learning of *common background knowledge* that can be used for future learning, reasoning and planning. The work at Carnegie Mellon University on NELL is an early example of such research [Carlson *et al.*, 2010].

### 3.4 Effective and Efficient Knowledge Transfer

The search for transfer learning methods that are able to develop accurate (effective) hypotheses rapidly (efficient) is a challenging problem for both LML and NSI. Transfer learning should produce a hypothesis for a new task that meets or exceeds the generalization performance of a hypothesis developed from only the training examples. There is evidence that functional transfer (e.g. using secondary task examples and multiple task learning) surpasses that of representation transfer in its ability to produce more accurate hypotheses [Caruana, 1997; Silver and Poirier, 2004]. Conversely, research has shown that a representational form of knowledge transfer (initializing a neural network to the weights of a related task) is typically more efficient than a functional form but it rarely results in improved model effectiveness [Silver and Poirier, 2004].

### 3.5 Scalability

Scalability is often the most difficult and important challenges for artificial intelligent systems. For NSI systems the extraction of symbolic knowledge is normally demanding in terms of time complexity. In LML systems the processes of retention and transfer adds time and space complexity to the challenge of learning. A NSI or LML 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 a system grows linearly in all of these factors. The move to Big Data and commercial data analytics is applying pressure on artificial intelligence approaches such as NSI and LML.

### 3.6 Heterogenous Domains of Tasks

Although, much of NSI and LML research has focused on retention and transfer within a single domain of tasks, an important area of research will be the development of systems that work across heterogenous domains [Yang *et al.*, 2009]. In heterogeneous transfer learning, the key idea is to leverage the feature correspondence across heterogenous domains (such as images and tags; music and lyrics) to build an effective feature mapping for transferring knowledge. Having knowledge in symbolic form would provide a number of new avenues to pursue in terms of knowledge adaptation prior to transfer from one problem domain to another.

### 3.7 Acquisition and Use of Meta-knowledge

Both NSI and LML systems need to collect and retain meta-knowledge of their task domains. For example, it may be critical for a NSI system to retain knowledge of the relationship or relatedness between learned tasks and an LML system may need to save the range and resolution of its input attributes [Silver *et al.*, 2008].

### 3.8 Shared Application Domains

Software agents and robots have provided useful test platforms for empirical studies of NSI and LML systems [Thrun, 1996a]. Agents and robots will naturally need to learn new but related tasks. This will provide opportunities to try different methods of retaining and consolidating task knowledge. The agent's fixed input and output domains provide an environment to test the impact of curriculum and the practice of tasks in a controlled manner. NSI and LML can also be used to overcome the *cold-start* problem exhibited by personal agents that employ user modeling [Lashkari *et al.*, 1994]. Retained knowledge can be used to boot-strap a new user model by transferring knowledge from a related user model.

## 4 Differences in NSI and LML Objectives

The following are some of the fundamental differences that NSI and LML researchers may have in their motives and objectives. These need to be appreciated by both research groups as they move to work together.

### 4.1 Uses of Retained Common Knowledge

LML systems focus on the use of learned knowledge for deployment in software systems or transfer during future learning. There is not as much consideration given to the transparency of learned knowledge, its use for explanation, or refinement of symbolic knowledge. In NSI systems the refinement and improvement of knowledge during learning is very important. Similarly, the ability to extract symbolic logic in a human readable form is important.

### 4.2 Retention versus Consolidation

LML researchers have started to consider the structure of retained knowledge for a lifelong learning system. Knowledge retention is necessary, but it may not be sufficient. In [Silver and Poirier, 2004] we propose that domain knowledge must be integrated for the purposes of efficient and effective retention and for more efficient and effective transfer during future learning. The process of integration we define as *consolidation*. The challenge for a LML system is consolidating the knowledge of a new task while retaining and possibly improving knowledge of prior tasks. An interesting aspect of this research is overcoming the *stability-plasticity* problem. The stability-plasticity problem refers to the challenge of adding new information to a system without the loss of prior information [Grossberg, 1987].

Traditional NSI systems tend to focus on methods of retaining symbolic representations of learned knowledge. Consolidation of prior knowledge with new knowledge is not often discussed. A recent survey by Lamb points out that connectionist methods of representing symbolic concepts is destroying NSI myths that have been around for some time [Lamb, 2008]. This may provide new ground for discovery of methods of consolidating knowledge in NSI systems.

### 4.3 Curriculum versus Expert Knowledge

An area where NSI and LML research differs is the use of expert knowledge. NSI researchers pride themselves on the ability to prime their learning systems with knowledge provided by an expert user. Such knowledge can be used as a source of inductive bias that leverages the available training examples. In contrast, most LML researchers pride themselves on designing self-taught leaners, and investigating curriculum and training sequences that are beneficial for learning a collection of increasingly complex tasks from as little knowledge as possible.

### 4.4 Practicing a Task

LML considers systems that may practice one or more tasks over a lifetime. The expectation is that the accuracy of the models of these tasks, residing in common knowledge, increases over time. This is closely related to the idea of knowledge consolidation versus simple retention. A computational theory of how best to practice tasks is important to artificial intelligence, as well as psychology and education. To the best of our knowledge, this is not an area that has been studied by NSI researchers, but their differing perspectives may lead to novel approaches.

### 4.5 Reasoning versus Learning

NSI systems are able to learn new knowledge, retain it in symbolic form, and then reason with the symbolic representation of the knowledge. They also consider the use of symbolic knowledge for transfer learning. LML systems tend to focus exclusively on knowledge retention for the purposes of transfer learning. LML researchers could benefit from a better understanding of the constraints placed on a learning system when the knowledge acquired must be amenable to reasoning. These constraints can be consider additional inductive biases that may be informative with respect to representation and search when learning.

## 5 Conclusion

In this paper we have proposed that research in Neural-Symbolic Integration and Lifelong Machine Learning share common ground that can be further exploited for the advancement of artificial intelligence. We have also pointed out several differences in NSI and LML research motives and objectives for further discussion by the community. We encourage researchers to explore these differences as interesting areas for making new discoveries in machine learning, knowledge representation and reasoning. In particular, we foresee that joint research has the potential to make serious advances in the areas of learning and use of common background knowledge that is so important for all areas of artificial intelligence.

A Dagstuhl seminar on Neural-Symbolic Learning and Reasoning is planned for September, 2014 [Garcez *et al.*, 2014]. The goal of the seminar is to build bridges between symbolic and sub-symbolic reasoning and learning representations using computer vision as a catalyst application. We see this as an opportunity for a number of artificial intelligence researchers to consider the links between LML and NSI.

## Acknowledgments

This research has been funded, in part, by the Natural Science and Engineering Research Council of Canada. Portions of this paper were extracted from recent joint work with Qiang Yang and Lianghao Li [Silver *et al.*, 2013].

## References

[Bader and Hitzler, 2005] Sebastian Bader and Pascal Hitzler. Dimensions of neural-symbolic integration - a structured survey. In *We Will Show Them: Essays in Honour of Dov Gabbay*, pages 167–194. College Publications, 2005.

[Bengio, 2009] Yoshua Bengio. Learning deep architectures for ai. *Foundations and Trends in Machine Learning*, 2(1):1–127, 2009.

[Carlson *et al.*, 2010] Andrew Carlson, Justin Betteridge, Bryan Kisiel, Burr Settles, Estevam R. Hruschka Jr., and Tom M. Mitchell. Toward an architecture for never-ending language learning. In Maria Fox and David Poole, editors, *AAAI*. AAAI Press, 2010.

[Caruana, 1997] Richard A. Caruana. Multitask learning. *Machine Learning*, 28:41–75, 1997.

[Garcez and Gabbay, 2004] Artur S. d'Avila Garcez and Dov M. Gabbay. Fibring neural networks. In Deborah L. McGuinness and George Ferguson, editors, *AAAI*, pages 342–347. AAAI Press / The MIT Press, 2004.

[Garcez *et al.*, 2002] Artur S. d'Avila Garcez, Dov M. Gabbay, and Krysia B. Broda. *Neural-Symbolic Learning System: Foundations and Applications*. Springer-Verlag New York, Inc., Secaucus, NJ, USA, 2002.

[Garcez *et al.*, 2009] Artur S. d'Avila Garcez, Luís C. Lamb, and Dov M. Gabbay. *Neural-Symbolic Cognitive Reasoning*. Cognitive Technologies. Springer, 2009.

[Garcez *et al.*, 2014] Artur Garcez, Marco Gori, Pascal Hitzler, and Luis Lamb. Dagstuhl seminar - neural-symbolic learning and reasoning, September 2014.

[Grossberg, 1987] Stephen Grossberg. Competitive learning: From interactive activation to adaptive resonance. *Cognitive Science*, 11:23–64, 1987.

[Lamb, 2008] Luís C. Lamb. The grand challenges and myths of neural-symbolic computation. In Luc De Raedt, Barbara Hammer, Pascal Hitzler, and Wolfgang Maass, editors, *Recurrent Neural Networks*, volume 08041 of *Dagstuhl Seminar Proceedings*. Internationales Begegnungs- und Forschungszentrum fuer Informatik (IBFI), Schloss Dagstuhl, Germany, 2008.

[Lashkari *et al.*, 1994] Yezdi Lashkari, Max Metral, and Pattie Maes. Collaborative interface agents. In *In Proceedings of the Twelfth National Conference on Artificial Intelligence*, pages 444–449. AAAI Press, 1994.

[Le *et al.*, 2012] Quoc Le, Marc'Aurelio Ranzato, Rajat Monga, Matthieu Devin, Kai Chen, Greg Corrado, Jeff Dean, and Andrew Ng. Building high-level features using large scale unsupervised learning. In *International Conference in Machine Learning*, 2012.

[Michalski, 1993] R.S. Michalski. Learning = inferencing + memorizing. *Foundations of Knowledge Acquistion: Machine Learning*, pages 1–41, 1993.

[Mitchell, 1980] Tom. M. Mitchell. The need for biases in learning generalizations. *Readings in Machine Learning*, pages 184–191, 1980. ed. Jude W. Shavlik and Thomas G. Dietterich.

[Raina *et al.*, 2007] Rajat Raina, Alexis Battle, Honglak Lee, Benjamin Packer, and Andrew Y. Ng. Self-taught learning: transfer learning from unlabeled data. In *Proceedings of the 24th international conference on Machine learning*, ICML '07, pages 759–766, New York, NY, USA, 2007. ACM.

[Ring, 1997] Mark B. Ring. Child: A first step towards continual learning. In *Machine Learning*, pages 77–104, 1997.

[Shavlik, 1992] Jude W. Shavlik. A framework for combining symbolic and neural learning. In *Machine Learning*, pages 321–331. Academic Press, 1992.

[Shultz and Rivest, 2001] Thomas R. Shultz and François Rivest. Knowledge-based cascade-correlation: using knowledge to speed learning. *Connect. Sci.*, 13(1):43–72, 2001.

[Silver and Mercer, 2002] Daniel L. Silver and Robert E. Mercer. The task rehearsal method of life-long learning: Overcoming impoverished data. *Advances in Artificial Intelligence, 15th Conference of the Canadian Society for Computational Studies of Intelligence (AI'2002)*, pages 90–101, 2002.

[Silver and Poirier, 2004] Daniel L. Silver and Ryan Poirier. Sequential consolidation of learned task knowledge. *Lecture Notes in Artificial Intelligence, 17th Conference of the Canadian Society for Computational Studies of Intelligence (AI'2004)*, pages 217–232, 2004.

[Silver *et al.*, 2008] Daniel L. Silver, Ryan Poirier, and Duane Currie. Inductive tranfser with context-sensitive neural networks. *Machine Learning*, 73(3):313–336, 2008.

[Silver *et al.*, 2013] Daniel L. Silver, Qiang Yang, and Lianghao Li. Lifelong machine learning systems: Beyond learning algorithms. In *Proceedings of the AAAI Spring Symposium on LifeLong Machine Learning, Stanford University, CA*, pages 49–55. AAAI, 2013.

[Solomonoff, 1989] Ray J. Solomonoff. 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*, pages 515–527, 1989.

[Sutton *et al.*, 2007] Richard S. Sutton, Anna Koop, and David Silver. On the role of tracking in stationary environments. In *Proceedings of the 24th international conference on Machine learning*, ICML '07, pages 871–878, New York, NY, USA, 2007. ACM.

[Tanaka and Yamamura, 1999] Fumihide Tanaka and Masayuki Yamamura. An approach to lifelong reinforcement learning through multiple environments. In *Proc. of the 6th European Workshop on Learning Robots (EWLR-6)*, pages 93–99, 1999.

[Thrun and O'Sullivan, 1995] Sebastian Thrun and J. O'Sullivan. Clustering learning tasks and the selective cross-task transfer of knowledge. *Technical Report CMU-CS-95-209*, Nov 1995.

[Thrun, 1996a] Sebastian Thrun. *Explanation-Based Neural Network Learning: A Lifelong Learning Approach*. Kluwer Academic Publishers, Boston, MA, 1996.

[Thrun, 1996b] Sebastian Thrun. Is learning the n-th thing any easier than learning the first? In *Advances in Neural Information Processing Systems*, pages 640–646. The MIT Press, 1996.

[Thrun, 1997] Sebastian Thrun. Lifelong learning algorithms. *Learning to Learn*, pages 181–209.1 Kluwer Academic Publisher, 1997.

[Tran and Garcez, 2012] S. Tran and A. S. d'Avila Garcez. Logic extraction from deep belief networks. In *Proceedings of ICML Workshop on Representation Learning*, Edinburgh, Scotland, 2012.

[Utgoff, 1983] Paul E. Utgoff. Adjusting bias in concept learning. In *Proceedings of IJCAI-1983*, pages 447–449, 1983.

[Wolpert, 1996] David H. Wolpert. The lack of a priori distinctions between learning algorithms. *Neural Comput.*, 8(7):1341–1390, October 1996.

[Yang *et al.*, 2009] Qiang Yang, Yuqiang Chen, Gui-Rong Xue, Wenyuan Dai, and Yong Yu. Heterogeneous transfer learning for image clustering via the social web. In *Proc. of the Joint Conf. of the 47th Annual Meeting of the ACL and the 4th Int. Joint Conf. on NLP of the AFNLP: Volume 1*, ACL '09, pages 1–9, Stroudsburg, PA, USA, 2009. Assoc. for Computational Linguistics.