Artificial Intelligence 254 (2018) 21–43

Contents lists available at ScienceDirect

Artificial Intelligence

www.elsevier.com/locate/artint

Overlapping layered learning ☆

Patrick MacAlpine *, Peter Stone *

Department of Computer Science, The University of Texas at Austin, Austin, TX 78701, USA

ARTICLE INFO

Article history:
Received 13 April 2017
Accepted 11 September 2017
Available online xxxx

Keywords:
Layered learning
Reinforcement learning
Robot soccer
CMA-ES
Robot skill learning
Hierarchical machine learning

ABSTRACT

Layered learning is a hierarchical machine learning paradigm that enables learning of complex behaviors by incrementally learning a series of sub-behaviors. A key feature of layered learning is that higher layers directly depend on the learned lower layers. In its original formulation, lower layers were frozen prior to learning higher layers. This article considers a major extension to the paradigm that allows learning certain behaviors independently, and then later stitching them together by learning at the "seams" where their influences overlap. The UT Austin Villa 2014 RoboCup 3D simulation team, using such overlapping layered learning, learned a total of 19 layered behaviors for a simulated soccer-playing robot, organized both in series and in parallel. To the best of our knowledge this is more than three times the number of layered behaviors in any prior layered learning system. Furthermore, the complete learning process is repeated on four additional robot body types, showcasing its generality as a paradigm for efficient behavior learning. The resulting team won the RoboCup 2014 championship with an undefeated record, scoring 52 goals and conceding none. This article includes a detailed experimental analysis of the team's performance and the overlapping layered learning approach that led to its success.

1. Introduction

Task decomposition is a popular approach for learning complex control tasks when monolithic learning—trying to learn the complete task all at once—is difficult or intractable [4–6]. Layered learning [7] is a hierarchical task decomposition machine learning paradigm that enables learning of complex behaviors by incrementally learning a series of sub-behaviors. A key feature of layered learning is that higher layers directly depend on the learned lower layers. In its original formulation, lower layers were frozen prior to learning higher layers. Freezing lower layers can be restrictive, however, as doing so limits the combined behavior search space over all layers. Concurrent layered learning [8] reduced this restriction in the search space by introducing the possibility of learning some of the behaviors simultaneously by "reopening" learning at the lower layers while learning the higher layers. A potential drawback of increasing the size of the search space, however, is an increase in the dimensionality and thus possibly the difficulty of what is being learned.

This article considers an extension to the layered learning paradigm, known as *overlapping layered learning*, that allows learning certain parameterized behaviors independently, and then later stitching them together by learning at the "seams" where their influences overlap. Overlapping layered learning aims to provide a middle ground between reductions in the

☆ This article is based on previous conference publications [1–3].
* Corresponding authors.
E-mail addresses: patmac@cs.utexas.edu (P. MacAlpine), pstone@cs.utexas.edu (P. Stone).

https://doi.org/10.1016/j.artint.2017.09.001
0004-3702/

Table 1
The key principles of layered learning.

1. A mapping directly from inputs to outputs is not tractably learnable
2. A bottom–up, hierarchical task decomposition is given
3. Machine learning exploits data to train and/or adapt. Learning occurs separately at each level
4. The output of learning in one layer feeds into the next layer

search space caused by freezing previously learned layers and the increased dimensionality of concurrent layered learning. Additionally, for complex tasks where it is difficult to learn one subtask in the presence of another, it reduces the dimensionality of the parameter search space by focusing only on parts responsible for subtasks working together.

The UT Austin Villa 2014 RoboCup 3D simulation team, using overlapping layered learning, learned a total of 19 layered behaviors for a simulated soccer-playing robot, organized both in series and in parallel. To the best of our knowledge this is more than three times the number of layered behaviors in any prior layered learning system. Furthermore, the complete learning process is repeated on four different heterogeneous robot body types, showcasing its generality as a paradigm for efficient behavior learning. The resulting team won the RoboCup 2014 championship with an undefeated record, scoring 52 goals and conceding none.

Primary contributions of this article are twofold. First, we introduce the overlapping layered learning paradigm, and present general scenarios where its use is beneficial. Second, we provide a detailed description and analysis of our machine learning approach, incorporating overlapping layered learning, to create a large and complex control system that was a core component of the 2014 RoboCup 3D simulation league championship team as well as three subsequent championship teams.

The remainder of this article is organized as follows. Section 2 provides background information on the original layered learning paradigm which is the basis for this work. Section 3 specifies and motivates the overlapping layered learning paradigm while contrasting it with traditional and concurrent layered learning. In Section 4 we introduce the RoboCup 3D simulation domain in which we evaluate this research. Section 5 details the overlapping layered learning approach of the 2014 UT Austin Villa team and in Section 6 we provide detailed analysis of its performance. Section 7 discusses related work while Section 8 concludes.

2. Layered learning paradigm

Table 1 summarizes the principles of the original layered learning paradigm which are described in detail in this section.[1]

2.1. Principle 1

Layered learning is designed for domains that are too complex for learning a mapping directly from the input to the output representation. Instead, the layered learning approach consists of breaking a problem down into several task layers. At each layer, a concept needs to be acquired. A machine learning (ML) algorithm abstracts and solves the local concept-learning task.

2.2. Principle 2

Layered learning uses a bottom–up incremental approach to hierarchical task decomposition. Starting with low-level subtasks, the process of creating new ML subtasks continues until reaching the high-level task that deal with the full domain complexity. The appropriate learning granularity and subtasks to be learned are determined as a function of the specific domain. The task decomposition in layered learning is not automated. Instead, the layers are defined by the ML opportunities in the domain.

2.3. Principle 3

Machine learning is used as a central part of layered learning to exploit data in order to train and/or adapt the overall system. ML is useful for training functions that are difficult to fine-tune manually. It is useful for adaptation when the task details are not completely known in advance or when they may change dynamically. In the former case, learning can be done off-line and frozen for future use. In the latter, on-line learning is necessary: since the learner needs to adapt to unexpected situations, it must be able to alter its behavior even while executing its task. Like the task decomposition itself, the choice of machine learning method depends on the subtask.

[1] This section is adapted from [7].