Transfer Learning for Multiagent Reinforcement Learning Systems

Synthesis Lectures on Artificial Intelligence and Machine Learning

Series Page

Transfer Learning for Multiagent Reinforcement Learning Systems
Felipe Leno da Silva and Anna Helena Reali Costa
2021

Network Embedding: Theories, Methods, and Applications
Cheng Yang, Zhiyuan Liu, Cunchao Tu, Chuan Shi, and Maosong Sun
2021

Introduction to Symbolic Plan and Goal Recognition
Reuth Mirsky, Sarah Keren, and Christopher Geib
2021

Graph Representation Learning
William L. Hamilton
2020

Introduction to Graph Neural Networks
Zhiyuan Liu and Jie Zhou
2020

Introduction to Logic Programming
Michael Genesereth and Vinay Chaudhri
2020

Federated Learning
Qiang Yang, Yang Liu, Yong Cheng, Yan Kang, and Tianjian Chen
2019

An Introduction to the Planning Domain Definition Language
Patrik Haslum, Nir Lipovetzky, Daniele Magazzeni, and Christina Muise
2019

Reasoning with Probabilistic and Deterministic Graphical Models: Exact Algorithms, Second Edition
Rina Dechter
2019

Learning and Decision-Making from Rank Data
Liron Xia
2019

Lifelong Machine Learning, Second Edition
Zhiyuan Chen and Bing Liu
2018

Adversarial Machine Learning
Yevgeniy Vorobeychik and Murat Kantarcioglu
2018

Strategic Voting
Reshef Meir
2018

Predicting Human Decision-Making: From Prediction to Action
Ariel Rosenfeld and Sarit Kraus
2018

Game Theory for Data Science: Eliciting Truthful Information
Boi Faltings and Goran Radanovic
2017

Multi-Objective Decision Making
Diederik M. Roijers and Shimon Whiteson
2017

Lifelong Machine Learning
Zhiyuan Chen and Bing Liu
2016

Statistical Relational Artificial Intelligence: Logic, Probability, and Computation
Luc De Raedt, Kristian Kersting, Sriraam Natarajan, and David Poole
2016

Representing and Reasoning with Qualitative Preferences: Tools and Applications
Ganesh Ram Santhanam, Samik Basu, and Vasant Honavar
2016

Metric Learning
Aurélien Bellet, Amaury Habrard, and Marc Sebban
2015

Graph-Based Semi-Supervised Learning
Amarnag Subramanya and Partha Pratim Talukdar
2014

Robot Learning from Human Teachers
Sonia Chernova and Andrea L. Thomaz
2014

General Game Playing
Michael Genesereth and Michael Thielscher
2014

Judgment Aggregation: A Primer
Davide Grossi and Gabriella Pigozzi
2014

An Introduction to Constraint-Based Temporal Reasoning
Roman Barták, Robert A. Morris, and K. Brent Venable
2014

Reasoning with Probabilistic and Deterministic Graphical Models: Exact Algorithms
Rina Dechter
2013

Introduction to Intelligent Systems in Traffic and Transportation
Ana L.C. Bazzan and Franziska Klügl
2013

A Concise Introduction to Models and Methods for Automated Planning
Hector Geffner and Blai Bonet
2013

Essential Principles for Autonomous Robotics
Henry Hexmoor
2013

Case-Based Reasoning: A Concise Introduction
Beatriz López
2013

Answer Set Solving in Practice
Martin Gebser, Roland Kaminski, Benjamin Kaufmann, and Torsten Schaub
2012

Planning with Markov Decision Processes: An AI Perspective
Mausam and Andrey Kolobov
2012

Active Learning
Burr Settles
2012

Computational Aspects of Cooperative Game Theory
Georgios Chalkiadakis, Edith Elkind, and Michael Wooldridge
2011

Representations and Techniques for 3D Object Recognition and Scene Interpretation
Derek Hoiem and Silvio Savarese
2011

A Short Introduction to Preferences: Between Artificial Intelligence and Social Choice
Francesca Rossi, Kristen Brent Venable, and Toby Walsh
2011

Human Computation
Edith Law and Luis von Ahn
2011

Trading Agents
Michael P. Wellman
2011

Visual Object Recognition
Kristen Grauman and Bastian Leibe
2011

Learning with Support Vector Machines
Colin Campbell and Yiming Ying
2011

Algorithms for Reinforcement Learning
Csaba Szepesvári
2010

Data Integration: The Relational Logic Approach
Michael Genesereth
2010

Markov Logic: An Interface Layer for Artificial Intelligence
Pedro Domingos and Daniel Lowd
2009

Introduction to Semi-Supervised Learning
Xiaojin Zhu and Andrew B.Goldberg
2009

Action Programming Languages
Michael Thielscher
2008

Representation Discovery using Harmonic Analysis
Sridhar Mahadevan
2008

Essentials of Game Theory: A Concise Multidisciplinary Introduction
Kevin Leyton-Brown and Yoav Shoham
2008

A Concise Introduction to Multiagent Systems and Distributed Artificial Intelligence
Nikos Vlassis
2007

Intelligent Autonomous Robotics: A Robot Soccer Case Study
Peter Stone
2007

Transfer Learning for Multiagent Reinforcement Learning Systems

Felipe Leno da Silva and Anna Helena Reali Costa

ISBN: 978-3-031-00463-6 paperback
ISBN: 978-3-031-01591-5 ebook
ISBN: 978-3-031-00036-2 hardcover

DOI 10.1007/978-3-031-01591-5

A Publication in the Springer series
SYNTHESIS LECTURES ON ARTIFICIAL INTELLIGENCE AND MACHINE LEARNING

Lecture #49
Series Editors: Ronald Brachman, *Jacobs Technion-Cornell Institute at Cornell Tech*
Francesca Rossi, *IBM Research AI*
Peter Stone, *University of Texas at Austin*
Series ISSN
Print 1939-4608 Electronic 1939-4616

Transfer Learning for Multiagent Reinforcement Learning Systems

Felipe Leno da Silva
Advanced Institute for AI

Anna Helena Reali Costa
Universidade de São Paulo

SYNTHESIS LECTURES ON ARTIFICIAL INTELLIGENCE AND MACHINE LEARNING #49

ABSTRACT

Learning to solve sequential decision-making tasks is difficult. Humans take years exploring the environment essentially in a random way until they are able to reason, solve difficult tasks, and collaborate with other humans towards a common goal. Artificial Intelligent agents are like humans in this aspect. Reinforcement Learning (RL) is a well-known technique to train autonomous agents through interactions with the environment. Unfortunately, the learning process has a high sample complexity to infer an effective actuation policy, especially when multiple agents are simultaneously actuating in the environment.

However, previous knowledge can be leveraged to accelerate learning and enable solving harder tasks. In the same way humans build skills and reuse them by relating different tasks, RL agents might reuse knowledge from previously solved tasks and from the exchange of knowledge with other agents in the environment. In fact, virtually all of the most challenging tasks currently solved by RL rely on embedded knowledge reuse techniques, such as Imitation Learning, Learning from Demonstration, and Curriculum Learning.

This book surveys the literature on knowledge reuse in multiagent RL. The authors define a unifying taxonomy of state-of-the-art solutions for reusing knowledge, providing a comprehensive discussion of recent progress in the area. In this book, readers will find a comprehensive discussion of the many ways in which knowledge can be reused in multiagent sequential decision-making tasks, as well as in which scenarios each of the approaches is more efficient. The authors also provide their view of the current low-hanging fruit developments of the area, as well as the still-open big questions that could result in breakthrough developments. Finally, the book provides resources to researchers who intend to join this area or leverage those techniques, including a list of conferences, journals, and implementation tools.

This book will be useful for a wide audience; and will hopefully promote new dialogues across communities and novel developments in the area.

KEYWORDS

multiagent reinforcement learning, transfer learning, learning from demonstrations, imitation learning, multi-task learning, knowledge reuse, machine learning, artificial intelligence

Contents

Preface

Artificial Intelligence (AI) systems are becoming so pervasive and integrated into our routine that adaptive behavior is rapidly becoming a necessity rather than an innovation in applications. *Learning* arose as the paradigm that brought AI to the spotlight, thanks to its ability to solve challenging tasks with minimal modeling effort. Personal virtual assistants, AI-powered robots, recommendation systems, smart devices, and autonomous cars were all science fiction material a couple of decades (perhaps years) ago, but are now part of our reality.

Despite recent successes, AI systems are still limited in reasoning about the long-term effects of their actions. *Reinforcement Learning* (RL) focuses on sequential decision-making problems. These techniques are based on an extensive exploration of the environment for learning action effects, which means that these techniques have a high sample complexity.

We firmly believe that collective efforts by artificial agents offer the path to crack yet-unsolved application challenges. Therefore, efficient techniques for training groups of agents to solve sequential decision-making tasks might be the next major AI breakthrough. However, in addition to the well-known high sample complexity that is amplified by the presence of multiple agents, multiagent RL has to cope with additional challenges such as the nonstationarity of other agents' behavior.

With this vision in mind, we have dedicated a good part of our scientific careers working on approaches for scaling up and facilitating the training of multiagent RL systems. As it happens with humans, reusing knowledge is a very effective way to accelerate the learning process. Hence, our approach has been to reuse knowledge to learn faster and to enable learning challenging tasks previously unsolvable through direct exploration.

The content that eventually became this book started to be written down years ago, as personal notes about the field, during the development of the first author's Ph.D. However, the material started to grow and to integrate works from different sub-communities such as *Supervised Learning*, *Multiagent Systems*, *Active Learning*, and of course *Reinforcement Learning*, each of them with their own (sometimes inconsistent) terminologies, jargon, and symbols. It was clear to us that an integrating effort was needed to systematically convey the content of the area.

In this context, we published a survey on *Transfer Learning in Multiagent Reinforcement Learning* in the *Journal of Artificial Intelligence Research* (JAIR). This book was written by extending that survey in several dimensions. First, although the survey was written only a couple of years ago, the area has had many developments after that, and we took this opportunity to update our picture of the state-of-the-art in the area. Second, the book format allowed us to reformulate the manuscript in a more didactic format. We included informative illustrations and described in more detail the background needed to fully understand the proposals in the area. We also updated our vision for the future based on the recent developments in the area.

This book was written to be accessible to practitioners and upper-level undergraduate students. Yet, it has enough in-deep content to be useful for Ph.D. students looking for a topic of research. We expect that this book will be useful for students, researchers, and practitioners interested in Transfer Learning, Reinforcement Learning, Multiagent Systems, and related areas. More than an interesting read, we expect that this book will inspire your research, and help you to see intersections with other communities you would not see otherwise.

Felipe Leno da Silva and Anna Helena Reali Costa
April 2021

Acknowledgments

We wish to acknowledge the collaboration of many colleagues with respect to this book. Be it from conference room discussions, university restaurant chats, or more formal paper reviews, too many people to be explicitly cited had an influence (sometimes anonymously) in this present content. F. L. Silva is especially grateful to Matthew E. Taylor, Peter Stone, and Ruben Glatt for all the career and technical advice, as well as the many hours spent discussing research over the past few years. A. H. R. Costa would like to thank all her graduate students, current and former, for their valuable discussions and contributions over the years.

We would also like to thank Michael Morgan, Christine Kiilerich, Matt Taylor, the anonymous reviewers, and the entire Morgan & Claypool Publishers team involved in the review and editorial process of this book. Without their involvement, this book would probably never go out to the world.

We dedicate this book to our families, who laid the foundation of our education and supported us in many ways over the last years. *Obrigado Paula, Berto e Solange.* A. H. R. Costa would like to thanks her husband Fabio, her mother Rachel, and her children, Regina and Marina. They have helped in so many ways.

Finally, A. H. R. Costa gratefully acknowledges the partial funding from the Brazilian agencies *Brazilian National Council for Scientific and Technological Development* (CNPq) and *São Paulo Research Foundation* (FAPESP). F. L. Silva acknowledges partial funding from *Fundação para o Desenvolvimento da UNESP* (FUNDUNESP - 3061/2019-CCP).

Felipe Leno da Silva and Anna Helena Reali Costa
April 2021