

Lifelong Machine Learning

Second Edition

Synthesis Lectures on Artificial Intelligence and Machine Learning

Editors

Ronald J. Brachman, *Jacobs Technion-Cornell Institute at Cornell Tech*

Peter Stone, *University of Texas at Austin*

Lifelong Machine Learning, Second Edition

Zhiyuan Chen and Bing Liu

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

XiaojinZhu 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

© Springer Nature Switzerland AG 2022

Reprint of original edition © Morgan & Claypool 2018

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means—electronic, mechanical, photocopy, recording, or any other except for brief quotations in printed reviews, without the prior permission of the publisher.

Lifelong Machine Learning, Second Edition

Zhiyuan Chen and Bing Liu

ISBN: 978-3-031-00453-7 paperback

ISBN: 978-3-031-01581-6 ebook

ISBN: 978-3-031-02709-3 epub

ISBN: 978-3-031-00026-3 hardcover

DOI 10.1007/978-3-031-01581-6

A Publication in the Springer series

SYNTHESIS LECTURES ON ARTIFICIAL INTELLIGENCE AND MACHINE LEARNING

Lecture #38

Series Editors: Ronald J. Brachman, *Jacobs Technion-Cornell Institute at Cornell Tech*

Peter Stone, *University of Texas at Austin*

Series ISSN

Print 1939-4608 Electronic 1939-4616

Lifelong Machine Learning

Second Edition

Zhiyuan Chen

Google, Inc.

Bing Liu

University of Illinois at Chicago

*SYNTHESIS LECTURES ON ARTIFICIAL INTELLIGENCE AND
MACHINE LEARNING #38*

ABSTRACT

Lifelong Machine Learning, Second Edition is an introduction to an advanced machine learning paradigm that continuously learns by accumulating past knowledge that it then uses in future learning and problem solving. In contrast, the current dominant machine learning paradigm learns in isolation: given a training dataset, it runs a machine learning algorithm on the dataset to produce a model that is then used in its intended application. It makes no attempt to retain the learned knowledge and use it in subsequent learning. Unlike this isolated system, humans learn effectively with only a few examples precisely because our learning is very knowledge-driven: the knowledge learned in the past helps us learn new things with little data or effort. Lifelong learning aims to emulate this capability, because without it, an AI system cannot be considered truly intelligent.

Research in lifelong learning has developed significantly in the relatively short time since the first edition of this book was published. The purpose of this second edition is to expand the definition of lifelong learning, update the content of several chapters, and add a new chapter about continual learning in deep neural networks—which has been actively researched over the past two or three years. A few chapters have also been reorganized to make each of them more coherent for the reader. Moreover, the authors want to propose a unified framework for the research area. Currently, there are several research topics in machine learning that are closely related to lifelong learning—most notably, multi-task learning, transfer learning, and meta-learning—because they also employ the idea of knowledge sharing and transfer. This book brings all these topics under one roof and discusses their similarities and differences. Its goal is to introduce this emerging machine learning paradigm and present a comprehensive survey and review of the important research results and latest ideas in the area. This book is thus suitable for students, researchers, and practitioners who are interested in machine learning, data mining, natural language processing, or pattern recognition. Lecturers can readily use the book for courses in any of these related fields.

KEYWORDS

lifelong machine learning; lifelong learning; continuous learning; continual learning; meta-learning; never-ending learning; multi-task learning; transfer learning

Zhiyuan dedicates this book to his wife, Vena Li, and his parents.

*Bing dedicates this book to his wife, Yue He; his children,
Shelley and Kate; and his parents.*

Contents

Preface	xvii
Acknowledgments	xix
1 Introduction	1
1.1 Classic Machine Learning Paradigm	1
1.2 Motivating Examples	3
1.3 A Brief History of Lifelong Learning	6
1.4 Definition of Lifelong Learning	9
1.5 Types of Knowledge and Key Challenges	14
1.6 Evaluation Methodology and Role of Big Data	16
1.7 Outline of the Book	18
2 Related Learning Paradigms	21
2.1 Transfer Learning	21
2.1.1 Structural Correspondence Learning	22
2.1.2 Naïve Bayes Transfer Classifier	23
2.1.3 Deep Learning in Transfer Learning	24
2.1.4 Difference from Lifelong Learning	25
2.2 Multi-Task Learning	26
2.2.1 Task Relatedness in Multi-Task Learning	26
2.2.2 GO-MTL: Multi-Task Learning using Latent Basis	27
2.2.3 Deep Learning in Multi-Task Learning	29
2.2.4 Difference from Lifelong Learning	30
2.3 Online Learning	31
2.3.1 Difference from Lifelong Learning	31
2.4 Reinforcement Learning	32
2.4.1 Difference from Lifelong Learning	33
2.5 Meta Learning	33
2.5.1 Difference from Lifelong Learning	34
2.6 Summary	34

3	Lifelong Supervised Learning	35
3.1	Definition and Overview	36
3.2	Lifelong Memory-Based Learning	37
3.2.1	Two Memory-Based Learning Methods	37
3.2.2	Learning a New Representation for Lifelong Learning	37
3.3	Lifelong Neural Networks	38
3.3.1	MTL Net	38
3.3.2	Lifelong EBNN	39
3.4	ELLA: An Efficient Lifelong Learning Algorithm	40
3.4.1	Problem Setting	41
3.4.2	Objective Function	41
3.4.3	Dealing with the First Inefficiency	42
3.4.4	Dealing with the Second Inefficiency	44
3.4.5	Active Task Selection	45
3.5	Lifelong Naive Bayesian Classification	46
3.5.1	Naïve Bayesian Text Classification	46
3.5.2	Basic Ideas of LSC	48
3.5.3	LSC Technique	49
3.5.4	Discussions	50
3.6	Domain Word Embedding via Meta-Learning	51
3.7	Summary and Evaluation Datasets	53
4	Continual Learning and Catastrophic Forgetting	55
4.1	Catastrophic Forgetting	55
4.2	Continual Learning in Neural Networks	57
4.3	Learning without Forgetting	59
4.4	Progressive Neural Networks	61
4.5	Elastic Weight Consolidation	62
4.6	iCaRL: Incremental Classifier and Representation Learning	64
4.6.1	Incremental Training	64
4.6.2	Updating Representation	65
4.6.3	Constructing Exemplar Sets for New Classes	66
4.6.4	Performing Classification in iCaRL	67
4.7	Expert Gate	67
4.7.1	Autoencoder Gate	68
4.7.2	Measuring Task Relatedness for Training	69

4.7.3	Selecting the Most Relevant Expert for Testing	69
4.7.4	Encoder-Based Lifelong Learning	70
4.8	Continual Learning with Generative Replay	70
4.8.1	Generative Adversarial Networks	70
4.8.2	Generative Replay	71
4.9	Evaluating Catastrophic Forgetting	72
4.10	Summary and Evaluation Datasets	73
5	Open-World Learning	77
5.1	Problem Definition and Applications	78
5.2	Center-Based Similarity Space Learning	79
5.2.1	Incrementally Updating a CBS Learning Model	79
5.2.2	Testing a CBS Learning Model	81
5.2.3	CBS Learning for Unseen Class Detection	82
5.3	DOC: Deep Open Classification	85
5.3.1	Feed-Forward Layers and the 1-vs.-Rest Layer	85
5.3.2	Reducing Open-Space Risk	86
5.3.3	DOC for Image Classification	88
5.3.4	Unseen Class Discovery	88
5.4	Summary and Evaluation Datasets	89
6	Lifelong Topic Modeling	91
6.1	Main Ideas of Lifelong Topic Modeling	91
6.2	LTM: A Lifelong Topic Model	94
6.2.1	LTM Model	95
6.2.2	Topic Knowledge Mining	96
6.2.3	Incorporating Past Knowledge	97
6.2.4	Conditional Distribution of Gibbs Sampler	99
6.3	AMC: A Lifelong Topic Model for Small Data	100
6.3.1	Overall Algorithm of AMC	100
6.3.2	Mining Must-link Knowledge	101
6.3.3	Mining Cannot-link Knowledge	103
6.3.4	Extended Pólya Urn Model	104
6.3.5	Sampling Distributions in Gibbs Sampler	106
6.4	Summary and Evaluation Datasets	108

7	Lifelong Information Extraction	111
7.1	NELL: A Never-Ending Language Learner	111
7.1.1	NELL Architecture	114
7.1.2	Extractors and Learning in NELL	114
7.1.3	Coupling Constraints in NELL	117
7.2	Lifelong Opinion Target Extraction	117
7.2.1	Lifelong Learning through Recommendation	118
7.2.2	AER Algorithm	119
7.2.3	Knowledge Learning	120
7.2.4	Recommendation using Past Knowledge	121
7.3	Learning on the Job	123
7.3.1	Conditional Random Fields	123
7.3.2	General Dependency Feature	124
7.3.3	The L-CRF Algorithm	126
7.4	Lifelong-RL: Lifelong Relaxation Labeling	127
7.4.1	Relaxation Labeling	127
7.4.2	Lifelong Relaxation Labeling	128
7.5	Summary and Evaluation Datasets	129
8	Continuous Knowledge Learning in Chatbots	131
8.1	LiLi: Lifelong Interactive Learning and Inference	132
8.2	Basic Ideas of LiLi	134
8.3	Components of LiLi	136
8.4	A Running Example	137
8.5	Summary and Evaluation Datasets	138
9	Lifelong Reinforcement Learning	139
9.1	Lifelong Reinforcement Learning through Multiple Environments	141
9.1.1	Acquiring and Incorporating Bias	141
9.2	Hierarchical Bayesian Lifelong Reinforcement Learning	142
9.2.1	Motivation	142
9.2.2	Hierarchical Bayesian Approach	143
9.2.3	MTRL Algorithm	143
9.2.4	Updating Hierarchical Model Parameters	144
9.2.5	Sampling an MDP	146
9.3	PG-ELLA: Lifelong Policy Gradient Reinforcement Learning	146

9.3.1	Policy Gradient Reinforcement Learning	147
9.3.2	Policy Gradient Lifelong Learning Setting	148
9.3.3	Objective Function and Optimization	149
9.3.4	Safe Policy Search for Lifelong Learning.....	150
9.3.5	Cross-domain Lifelong Reinforcement Learning	151
9.4	Summary and Evaluation Datasets	152
10	Conclusion and Future Directions	153
	Bibliography	159
	Authors' Biographies	187

Preface

The purpose of writing this second edition is to extend the definition of lifelong learning, to update the content of several chapters, and to add a new chapter about *continual learning in deep neural networks*, which has been actively researched for the past two to three years. A few chapters are also reorganized to make each of them more coherent.

The project of writing this book started with a tutorial on *lifelong machine learning* that we gave at the 24th International Joint Conference on Artificial Intelligence (IJCAI) in 2015. At that time, we had worked on the topic for a while and published several papers in ICML, KDD, and ACL. When Morgan & Claypool Publishers contacted us about the possibility of developing a book on the topic, we were excited. We strongly believe that lifelong machine learning (or simply lifelong learning) is very important for the future of machine learning and artificial intelligence (AI). Note that lifelong learning is sometimes also called *continual learning* or *continuous learning* in the literature. Our original research interest in the topic stemmed from extensive application experiences in sentiment analysis (SA) in a start-up company several years ago. A typical SA project starts with a client who is interested in consumer opinions expressed in social media about their products or services and those of their competitors. There are two main analysis tasks that an SA system needs to do: (1) discover the entities (e.g., *iPhone*) and entity attributes/features (e.g., *battery life*) that people talked about in opinion documents such as online reviews and (2) determine whether the opinion about each entity or entity attribute is positive, negative, or neutral [Liu, 2012, 2015]. For example, from the sentence “*iPhone is really cool, but its battery life sucks,*” an SA system should discover that the author is (1) positive about *iPhone* and (2) negative about *iPhone’s battery life*.

After working on many projects in many domains (which are types of products or services) for clients, we realized that there is a great deal of sharing of information across domains and projects. As we see more and more, new things get fewer and fewer. It is easy to see that sentiment words and expressions (such as *good*, *bad*, *poor*, *terrible*, and *cost an arm and a leg*) are shared across domains. There is also a great deal of sharing of entities and attributes. For example, every product has the attribute of *price*, most electronic products have *battery*, and many of them also have *screen*. It is silly not to exploit such sharing to significantly improve SA to make it much more accurate than without using such sharing but only working on each project and its data in isolation. The classic machine learning paradigm learns exactly in isolation. Given a dataset, a learning algorithm runs on the data to produce a model. The algorithm has no memory and thus is unable to use the previously learned knowledge. In order to exploit knowledge sharing, an SA system has to retain and accumulate the knowledge learned in the past and use it to help future learning and problem solving, which is exactly what *lifelong learning* aims to do.

It is not hard to imagine that this sharing of information or knowledge across domains and tasks is generally true in every field. It is particularly obvious in natural language processing because the meanings of words and phrases are basically the same across domains and tasks and so is the sentence syntax. No matter what subject matter we talk about, we use the same language, although each subject may use only a small subset of the words and phrases in a language. If that is not the case, it is doubtful that a natural language would have ever been developed by humans. Thus, lifelong learning is generally applicable, not just restricted to sentiment analysis.

The goal of this book is to introduce this emerging machine learning paradigm and to present a comprehensive survey and review of the important research results and latest ideas in the area. We also want to propose a unified framework for the research area. Currently, there are several research topics in machine learning that are closely related to lifelong learning, most notably, multi-task learning and transfer learning, because they also employ the idea of knowledge sharing and transfer. This book brings all these topics under one roof and discusses their similarities and differences. We see lifelong learning as an extension to these related paradigms. Through this book, we would also like to motivate and encourage researchers to work on lifelong learning. We believe it represents a major research direction for both machine learning and artificial intelligence for years to come. Without the capability of retaining and accumulating knowledge learned in the past, making inferences about it, and using the knowledge to help future learning and problem solving, achieving artificial general intelligence (AGI) is unlikely.

Two main principles have guided the writing of this book. First, it should contain strong motivations for conducting research in lifelong learning in order to encourage graduate students and researchers to work on lifelong learning problems. Second, the writing should be accessible to practitioners and upper-level undergraduate students who have basic knowledge of machine learning and data mining. Yet there should be sufficient in-depth materials for graduate students who plan to pursue Ph.D. degrees in the machine learning and/or data mining fields.

This book is thus suitable for students, researchers, and practitioners who are interested in machine learning, data mining, natural language processing, or pattern recognition. Lecturers can readily use the book in class for courses in any of these related fields.

Zhiyuan Chen and Bing Liu
August 2018

Acknowledgments

We would like to thank the current and former graduate students in our group and our collaborators: Geli Fei, Zhiqiang Gao, Estevam R. Hruschka Jr., Wenpeng Hu, Minlie Huang, Yongbing Huang, Doo Soon Kim, Huayi Li, Jian Li, Lifeng Liu, Qian Liu, Guangyi Lv, Sahisnu Mazumder, Arjun Mukherjee, Nianzu Ma, Lei Shu, Tao Huang, William Underwood, Hao Wang, Shuai Wang, Hu Xu, Yueshen Xu, Tim Yin, Tim Yuan, and Yuanlin Zhang, for their contributions of numerous research ideas and helpful discussions over the years. We are especially grateful to the two expert reviewers of the first edition, Eric Eaton and Matthew E. Taylor. Despite their busy schedules, they read the first draft of the book very carefully and gave us so many excellent comments and suggestions, which were not only insightful and comprehensive, but also detailed and very constructive. German I. Parisi reviewed Chapter 4 of this second edition and gave us many valuable comments. Their suggestions have helped us improve the book tremendously.

On the publication side, we thank the editors of Synthesis Lectures on Artificial Intelligence and Machine Learning, Ronald Brachman, William W. Cohen, and Peter Stone, for initiating this project. The President and CEO of Morgan & Claypool Publishers, Michael Morgan, and his staff, Christine Kiilerich, and C.L. Tondo have given us all kinds of help promptly whenever requested, for which we are very grateful.

Our greatest gratitude go to our own families. Zhiyuan Chen would like to thank his wife Vena Li and his parents. Bing Liu would like to thank his wife Yue, his children Shelley and Kate, and his parents. They have helped in so many ways.

The writing of this book was partially supported by two National Science Foundation (NSF) grants IIS-1407927 and IIS-1650900, an NCI grant R01CA192240, a research gift from Huawei Technologies, and a research gift from Robert Bosch GmbH. The content of the book is solely the responsibility of the authors and does not necessarily represent the official views of the NSF, NCI, Huawei, or Bosch. The Department of Computer Science at the University of Illinois at Chicago provided computing resources and a very supportive environment for this project. Working at Google has also given Zhiyuan Chen a broader perspective on machine learning.

Zhiyuan Chen and Bing Liu
August 2018