Learning and Decision-Making from Rank Data

Synthesis Lectures on Artificial Intelligence and Machine Learning

Editors
Ronald J. Brachman, ***Jacobs Technion-Cornell Institute at Cornell Tech***
Francesca Rossi, ***IBM Research AI***
Peter Stone, ***University of Texas at Austin***

Learning and Decision-Making from Rank Data
Lirong 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
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 2019

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.

Learning and Decision-Making from Rank Data

Lirong Xia

ISBN: 978-3-031-00454-4 paperback
ISBN: 978-3-031-01582-3 ebook
ISBN: 978-3-031-00027-0 hardcover

DOI 10.1007/978-3-031-01582-3

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

Lecture #40
Series Editors: Ronald J. 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

Learning and Decision-Making from Rank Data

Lirong Xia
Rensselaer Polytechnic Institute

SYNTHESIS LECTURES ON ARTIFICIAL INTELLIGENCE AND MACHINE LEARNING #40

ABSTRACT

The ubiquitous challenge of learning and decision-making from rank data arises in situations where intelligent systems collect preference and behavior data from humans, learn from the data, and then use the data to help humans make efficient, effective, and timely decisions. Often, such data are represented by *rankings*.

This book surveys some recent progress toward addressing the challenge from the considerations of statistics, computation, and socio-economics. We will cover classical statistical models for rank data, including random utility models, distance-based models, and mixture models. We will discuss and compare classical and state-of-the-art algorithms, such as algorithms based on Minorize-Majorization (MM), Expectation-Maximization (EM), Generalized Method-of-Moments (GMM), rank breaking, and tensor decomposition. We will also introduce principled Bayesian preference elicitation frameworks for collecting rank data. Finally, we will examine socio-economic aspects of statistically desirable decision-making mechanisms, such as Bayesian estimators.

This book can be useful in three ways: (1) for theoreticians in statistics and machine learning to better understand the considerations and caveats of learning from rank data, compared to learning from other types of data, especially cardinal data; (2) for practitioners to apply algorithms covered by the book for sampling, learning, and aggregation; and (3) as a textbook for graduate students or advanced undergraduate students to learn about the field.

This book requires that the reader has basic knowledge in probability, statistics, and algorithms. Knowledge in social choice would also help but is not required.

KEYWORDS

rank data, decision-making, random utility models, Plackett-Luce model, distance-based models, Mallows' model, preference elicitation, social choice, fairness, Bayesian estimators

Contents

Preface . . . . . **xiii**

Acknowledgments . . . . . **xv**

1 Introduction **1**

1.1 The Research Problem 4

1.2 Overview of the Book 7

2 Statistical Models for Rank Data **9**

2.1 Basics of Statistical Modeling 9

2.1.1 Modeling Partial Orders as Events 12

2.2 Random Utility Models 13

2.2.1 The Plackett–Luce Model 15

2.2.2 Properties of Random Utility Models (RUMs) 17

2.2.3 Sampling from Random Utility Models 20

2.2.4 Connection to Discrete Choice Models 22

2.3 Distance-Based Models 26

2.3.1 Mallows' Model 26

2.3.2 Repeated Insertion Model: Efficient Sampling from Mallows 28

2.3.3 Condorcet's Model 30

2.4 Datasets and Model Fitness 30

2.5 Bibliographical Notes 31

3 Parameter Estimation Algorithms **33**

3.1 Algorithms for the Plackett–Luce Model 33

3.1.1 The Minorize-Maximization (MM) Algorithm 35

3.1.2 The Luce Spectral Ranking (LSR) Algorithm 36

3.1.3 Generalized Method-of-Moments (GMM) Algorithm 38

3.2 Algorithms for General Random Utility Models 40

3.2.1 The Expectation-Maximization (EM) Algorithm 41

3.2.2 EM for RUMs: Monte Carlo E-Step by Gibbs Sampling 42

3.2.3 EM for RUMs: M-Step 45

3.2.4 GMM for RUMs with Location Families 45
3.3 Algorithms for Distance-Based Models 48
3.4 Bibliographical Notes 50

4 The Rank-Breaking Framework 53
4.1 Rank-Breaking for Random Utility Models 55
4.1.1 Breaking + GMM for Plackett–Luce 56
4.1.2 Uniqueness of Outcome of Algorithm 4.11 57
4.1.3 Characterization of Consistent Breakings for Plackett–Luce 58
4.1.4 Computational and Statistical Efficiency of Algorithms for Plackett–Luce 60
4.1.5 Rank-Breaking for General Random Utility Models with Location Families 62
4.2 Rank-Breaking + Composite Marginal Likelihood (RBCML) 64
4.2.1 Weighted Breakings 64
4.2.2 Composite Marginal Likelihood Methods (CML) 65
4.2.3 The RBCML Framework 66
4.2.4 Consistency and Asymptotic Normality of RBCML 67
4.2.5 RBCML for Plackett–Luce 69
4.2.6 RBCML for RUMs with Location Families 71
4.2.7 The Adaptive RBCML Algorithm 73
4.2.8 Experiments 73
4.3 Bibliographical Notes 75

5 Mixture Models for Rank Data 77
5.1 Mixture Models 77
5.1.1 Identifiability of Mixture Models 80
5.1.2 An EM Algorithm for Learning Mixture Models 80
5.2 Learning Mixtures of Plackett–Luce 81
5.2.1 Algorithms for Mixtures of Plackett–Luce 86
5.3 Learning Mixtures of General RUMs with Location Families 87
5.4 Learning Mixtures of Mallows 89
5.5 Bibliographical Notes 90

6 Bayesian Preference Elicitation 93
6.1 The Bayesian Preference Elicitation Problem 94
6.1.1 Plackett–Luce Model with Features 96

6.1.2 Computing Expected Information Gain 96
6.2 Bayesian Preference Elicitation for Personal Choice 98
6.2.1 Information Criteria 98
6.2.2 Approximation Techniques for Personal Choice 99
6.3 Bayesian Preference Elicitation for Social Choice 101
6.3.1 Approximating Posterior Distributions 101
6.3.2 Ranked-Top-k Questions 103
6.3.3 Social Choice by Randomized Voting 104
6.4 Experimental Results 106
6.4.1 Estimating the Cost Function 106
6.4.2 Comparing Information Criteria 108
6.5 Bibliographical Notes 109

7 Socially Desirable Group Decision-Making from Rank Data 111
7.1 Statistical Decision-Theoretic Framework 112
7.1.1 Measuring Decision Mechanisms: Bayesian Loss and Frequentist Loss 113
7.1.2 Socio-Economic Criteria: Social Choice Axioms 116
7.2 Minimax Estimators in Neutral Frameworks 117
7.3 Socially Desirable Bayesian Estimators 119
7.3.1 An Impossibility Theorem on Strict Condorcet Criterion 119
7.3.2 Satisfiability of Other Axioms 119
7.4 An Automated Design Framework 121
7.4.1 Data Generation 122
7.4.2 Hypothesis Space 123
7.4.3 Optimization 124
7.5 Bibliographical Notes 125

8 Future Directions 127

Bibliography 131

Author's Biography 143

Preface

We are living in a world of data. Every day, the booming IT industry collects large amounts of user preference and behavioral data to make various decisions. In many cases, these data are in various rank formats instead of numerical values over alternatives—for example, how voters rank candidates, consumers choose one product over another, search engines rank webpages, etc.

Efficiently collecting and learning from rank data are important challenges to artificial intelligence and machine learning. Recently, there has been a growing literature on preference learning and decision-making that leverages (1) classical statistical modeling for rank data, such as discrete choice models and random utility models, and (2) modern machine learning and optimization techniques. In addition, as the data are often collected from humans and the decisions are made for humans, socio-economic aspects such as fairness are also important considerations.

Most of the book aims to answer the question of how we can design fast and accurate algorithms for learning from rank data in terms of widely studied and applied statistical models. After introducing random utility models and distance-based models in Chapter 2, we discuss several classical and state-of-the-art parameter estimation algorithms in Chapter 3. Chapter 4 covers a general framework for handling rank data called rank-breaking. Chapter 5 focuses on the theory and algorithms of mixture models for rank data.

How to efficiently elicit preferences from users is the main topic of Chapter 6, where we take a Bayesian experimental design and active learning approach by computing the most cost-effective queries. This chapter also covers some Bayesian approaches toward learning from partial orders. Preliminary experiments show that choice data are not as cost-effective as other types of rank data.

Chapter 7 covers a statistical decision-theoretic framework for social choice, under which tradeoffs among statistical properties, fairness, and computation are studied. We consider how to make accurate and socially desirable decisions. An impossibility theorem on fair Bayesian estimators is shown, and positive results on the satisfaction of many desirable properties are discussed. The chapter also briefly discusses an automated framework for designing decision mechanisms with desirable statistical and fairness properties.

The major goal of this book is to stimulate future research and applications by surveying recent progress in learning and decision-making from rank data as well as to provide a reference to classical definitions, properties, and algorithms for handling rank data.

Lirong Xia
January 2019

Acknowledgments

I would like to thank my coauthors Hossein Azari Soufiani, William Chen, Yiling Chen, Vincent Conitzer, David Hughes, Kevin Hwang, Jeffrey Kephart, Jerome Lang, Haoming Li, Chao Liao, Ao Liu, Zhenming Liu, Pinyan Lu, Nicholas Mattei, David C. Parkes, Peter Piech, Matthew Rognlie, Sujoy Sikdar, Hui Su, Rohit Vaish, Tristan Villamil, Bo Waggoner, Junming Wang, Qiong Wu, Chaonan Ye, and Zhibing Zhao, for great collaborations and inspirations over years on topics covered by or closely related to this book. I would also like to thank Reshef Meir and Shivani Agarwal for extremely useful comments and feedbacks.

I also deeply appreciate the guidance and help from my postdoctoral and graduate school advisors David C. Parkes, Vincent Conitzer, and Mingsheng Ying.

This work is partly supported by NSF #1453542 and ONR #N00014-17-1-2621.

Last but not least, I want to thank my parents, my sister, my wife, and my daughter for giving me a wonderful life.

Lirong Xia
January 2019