Introduction to Symbolic Plan and Goal Recognition

Synthesis Lectures on Artificial Intelligence and Machine Learning

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

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

ISBN: 978-3-031-00461-2 paperback
ISBN: 978-3-031-01589-2 ebook
ISBN: 978-3-031-00034-8 hardcover

DOI 10.1007/978-3-031-01589-2

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

Lecture #47
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
Synthesis Lectures on Artificial Intelligence and Machine Learning
Print 1939-4608 Electronic 1939-4616

Introduction to Symbolic Plan and Goal Recognition

Reuth Mirsky
University of Texas at Austin

Sarah Keren
Harvard University

Christopher Geib
SIFT

SYNTHESIS LECTURES ON ARTIFICIAL INTELLIGENCE AND MACHINE LEARNING #47

ABSTRACT

Plan recognition, activity recognition, and goal recognition all involve making inferences about other actors based on observations of their interactions with the environment and other agents. This synergistic area of research combines, unites, and makes use of techniques and research from a wide range of areas including user modeling, machine vision, automated planning, intelligent user interfaces, human-computer interaction, autonomous and multi-agent systems, natural language understanding, and machine learning. It plays a crucial role in a wide variety of applications including assistive technology, software assistants, computer and network security, human-robot collaboration, natural language processing, video games, and many more.

This wide range of applications and disciplines has produced a wealth of ideas, models, tools, and results in the recognition literature. However, it has also contributed to fragmentation in the field, with researchers publishing relevant results in a wide spectrum of journals and conferences.

This book seeks to address this fragmentation by providing a high-level introduction and historical overview of the plan and goal recognition literature. It provides a description of the core elements that comprise these recognition problems and practical advice for modeling them. In particular, we define and distinguish the different recognition tasks. We formalize the major approaches to modeling these problems using a single motivating example. Finally, we describe a number of state-of-the-art systems and their extensions, future challenges, and some potential applications.

KEYWORDS

plan recognition, goal recognition, activity recognition, behavior recognition, intent recognition, temporal pattern recognition, reasoning under uncertainty, human–AI collaboration, multi-agent systems, symbolic reasoning

To Ilya, for giving me something greater than research, and then letting me have both.

- Reuth Mirsky

To my husband Uri and to my Ph.D. advisors, Avigdor Gal and Erez Karpas. I couldn't have done this without your support.

- Sarah Keren

To Sabra. Shockingly, sometimes words actually do fail me.... R&C.

- Christopher Geib

Contents

Preface

This book is based on a tutorial presented by the authors at the 2019 AAAI conference in Hawaii. The wish to create the tutorial came up after the authors had been co-chairing the Plan Activity and Intent (PAIR) workshop at AAAI since 2017. The PAIR workshop, initially named Modeling Others from Observations (MOO), has been taking place at different venues since 2004. Every year, it attracts and brings together researchers from diverse backgrounds and disciplines.

Our experience as co-chairs helped us appreciate the increasing interest in plan and goal recognition. However, it also highlighted the lack of a shared vocabulary and terminology to connect the different lines of work. Our intention for the tutorial and this book is to bridge the gap that exists between different threads of research in the field. We do this by providing an overview of past and state-of-the-art model-based plan and goal recognition literature, by specifying a formalization of the elements of the problem, and by describing a set of practical tools for evaluating and investigating a new recognition problem.

The book is organized into five chapters. Chapter 1 is an introduction to plan and goal recognition and an overview of key past works in the area. Chapter 2 provides a unified recipe for defining a recognition problem, and provides guidelines for choosing an approach for a given recognition task. In Chapter 3, we formalize goal and plan recognition. Chapter 4 describes a variety of state-of-the-art approaches to recognition and suggests ways to extend existing symbolic plan and goal recognition tools. Finally, Chapter 5 highlights possible directions for future work. In all of these chapters we provide references an interested reader can use to continue the exploration of this research space.

We hope this book will enable and encourage researchers to read more widely past work and to build on its lessons to advance plan and goal recognition research. To paraphrase Sir Isaac Newton, our research field truly has been built on the shoulders of giants.

Reuth Mirsky, Sarah Keren, and Christopher Geib
January 2021

Acknowledgments

While the cover of this book lists three authors, no such work is ever completed without the assistance, encouragement, and support of a strictly non-enumerable number of supporters, colleagues, editors, mentors, friends, and family (genetic and chosen). We thank you all. Further, we would like to acknowledge all of the exceptional researchers whose work we reference here. We have worked to the best of our ability at synthesizing a diverse field of research, and if in that effort our characterization of your work has been imperfect, we apologize. We would also like to acknowledge all of the authors of papers at the PAIR and MOO series of workshops that has been such a rich and supportive research community over the years and contributed to and shaped so much of our thinking about these problems. Gal Kaminka the "founding father" of the first MOO workshop deserves special acknowledgment in this regard. We would also like to thank all of the people that encouraged us (before and after) to run the Plan Activity and Intent Recognition Tutorial at AAAI-19 that was the impetus behind this volume.

We would all like to thank Michael Morgan and the whole team at Morgan & Claypool for their patience and support, Peter Stone for suggesting we turn our tutorial into a book, and our reviewers Gita Sukthankar, Shlomo Zilberstein, and David Pynadath for their comments on an earlier draft of this book. It has benefited immeasurably from their efforts and any errors in it were definitely inserted after their commentary. We would like to acknowledge and thank our partners and families for their patience and support during this process. Writing a book is never an easy or fast process and it is not only the authors that pay the costs. Thank you.

More specifically, Reuth would like to thank Kobi Gal for introducing her to the world of PAIR; Chris Geib and Gal Kaminka for being the Gandalf and the Elrond of her first PAIR adventures; Sarah Keren, Mor Vered, and Maayan Shvo for exciting discussions about PAIR and beyond; and Peter Stone for his inspiring mentorship and support.

Sarah would like to thank Avigdor Gal and Erez Karpas, her Ph.D. advisors who helped formulate Goal Recognition Design and introduce it to the world. She also thanks Barbara J. Grosz, David C. Parkes, and Jeffrey S. Rosenschein for supporting her in the process of writing this book and for helping her find new and better ways to make her messages clear.

Finally, Chris would like to thank his co-authors Reuth and Sarah. It is in our discussions, conversations, arguments, and efforts to write down what we know that we clarify our own

knowledge. The benefits of having you as such insightful, challenging, and thought-provoking, colleagues has been immeasurable. Thank you.

Reuth Mirsky, Sarah Keren, and Christopher Geib
January 2021