Case Based Reasoning A Concise Introduction Synthesis Lectures On Artificial Intelligence And Machine Learning

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.

Social choice theory deals with aggregating the preferences of multiple individuals regarding several available alternatives, a situation colloquially known as voting. There are many different voting rules in use and even more in the literature, owing to the various considerations such an aggregation method should take into account. The analysis of voting scenarios become particularly challenging in the presence of strategic voters, that is, voters that misreport their true preferences in an attempt to obtain a more favorable outcome. In a world that is tightly connected by the Internet, where multiple groups with complex incentives make frequent joint decisions, the interest in strategic voting exceeds the scope of political science and is a focus of research in economics, game theory, sociology, mathematics, and computer science. The book has two parts. The first part asks "are there voting rules that are truthful?" in the sense that all voters have an incentive to report their true preferences. The seminal Gibbard-Satterthwaite theorem excludes the existence of such voting rules under certain requirements. From this starting point, we survey both extensions of the theorem and various conditions under which truthful voting is made possible (such as restricted preference domains). We also explore the connections with other problems of mechanism design such as locating a facility that serves multiple users. In the second part, we ask "what would be the outcome when voters do vote strategically?" rather than trying to prevent such behavior. We overview various game-theoretic models and equilibrium concepts from the literature, demonstrate how they apply to voting games, and discuss their implications on social welfare. We conclude with a brief survey of

empirical and experimental findings that could play a key role in future development of game theoretic voting models.

This new resource provides a coherent, intuitive, and theoretical foundation for the fusion and exploitation of traditional sensor data as well as text-based information. In addition to presenting a detailed discussion of base-level data fusion requirements, a variety of higher level exploitation algorithms are presented that perform fully automated relationship discovery, rank interest level of entities, and support context-sensitive behavior understanding (both static and dynamic context). This book identifies eight canonical fusion forms as well as twenty foundational fusion services to enable formal mapping between models and services. Normalization and representation processes for (hard) sensor data and (soft) semantic data are described as well as methods for combining hard and soft data. Included is a prototype fusion system developed to implement virtually all the presented applications in order to demonstrate the robustness and utility of the design principles presented in this resource. The prototype system presented supports a variety of user workflows and all the applications are fully integrated. There is extensive fusion system output for unclassified scenarios to permit the reader to fully understand all presented design principles. This book also presents context-sensitive fuzzy semantic spatial and temporal reasoning.

Machine learning and artificial intelligence (AI) are powerful tools that create predictive models, extract information, and help make complex decisions. They do this by examining an enormous quantity of labeled training data to find patterns too complex for human observation.

However, in many real-world applications, well-labeled data can be difficult, expensive, or even impossible to obtain. In some cases, such as when identifying rare objects like new archeological sites or secret enemy military facilities in satellite images, acquiring labels could require months of trained human observers at incredible expense. Other times, as when attempting to predict disease infection during a pandemic such as COVID-19, reliable true labels may be nearly impossible to obtain early on due to lack of testing equipment or other factors. In that scenario, identifying even a small amount of truly negative data may be impossible due to the high false negative rate of available tests. In such problems, it is possible to label a small subset of data as belonging to the class of interest though it is impractical to manually label all data not of interest. We are left with a small set of positive labeled data and a large set of unknown and unlabeled data. Readers will explore this Positive and Unlabeled learning (PU learning) problem in depth. The book rigorously defines the PU learning problem, discusses several common assumptions that are frequently made about the problem and their implications, and considers how to evaluate solutions for this problem before describing several of the most popular algorithms to solve this problem. It explores several uses for PU learning including applications in biological/medical, business, security, and signal processing. This book also provides high-level summaries of several related learning problems such as one-class classification, anomaly detection, and noisy learning and their relation to PU learning.

Eliciting Truthful Information

A Concise Introduction

4th International Conference on Case-Based Reasoning, ICCBR 2001 Vancouver, BC, Canada, July 30 - August 2, 2001 Proceedings

Network Embedding

Computer Science And Technology - Proceedings Of The International Conference (Cst2016)

Graph-structured data is ubiquitous throughout the natural and social sciences, from telecommunication networks to quantum chemistry. Building relational inductive biases into deep learning architectures is crucial for creating systems that can learn, reason, and generalize from this kind of data. Recent years have seen a surge in research on graph representation learning, including techniques for deep graph embeddings, generalizations of convolutional neural networks to graph-structured data, and neural message-passing approaches inspired by belief propagation. These advances in graph representation learning have led to new state-of-the-art results in numerous domains, including chemical synthesis, 3D vision, recommender systems, question answering, and social network analysis. This book provides a synthesis and overview of graph representation learning. It begins with a discussion of the goals of graph representation learning as well as key methodological foundations in graph theory and network analysis. Following this, the book introduces and reviews methods for learning node embeddings, including random-walk-based methods and applications to knowledge graphs. It then provides a technical synthesis and introduction to the highly successful graph neural network (GNN) formalism, which has become a dominant and fast-growing paradigm for deep learning with graph data. The book concludes with a synthesis of recent advancements in deep generative

models for graphs—a nascent but quickly growing subset of graph representation learning.
This proceedings consists of selected papers presented at the International Conference on Computer Science and Technology (CST2016), which was successfully held in Shenzhen, China during January 8–10, 2016. CST2016 covered a wide range of fundamental studies, technical innovations and industrial applications in 7 areas, namely Computer Systems, Computer Network, Security, Databases and Information Systems, Artificial Intelligence and Multimedia, Theory and Software Engineering and Computer Applications. CST 2016 aims to provide a forum for researchers, engineers, and students in the area of computer science and technology. It features unique mixed various topics in computer science and technology including big data, system architecture, hardware and applications. CST 2016 attracted more than 300 submissions. Among them, only 142 papers were accepted in to the conference after a stringent peer review process.

How are new technologies changing the practice of law? With examples and explanations drawn from the UK, US, Canada, Australia and other common law countries, as well as from China and Europe, this book considers the opportunities and implications for lawyers as artificial intelligence systems become commonplace in legal service delivery. It examines what lawyers do in the practice of law and where AI will impact this work. It also explains the important continuing role of the lawyer in an AI world. This book is divided into three parts: Part A provides an accessible explanation of AI, including diagrams, and contrasts this with the role and work of lawyers. Part B focuses on six different aspects of legal work (litigation, transactional,

dispute resolution, regulation and compliance, criminal law and legal advice and strategy) where AI is making a considerable impact and looks at how this is occurring. Part C discusses how lawyers and law firms can best utilise the promise of AI, while also acknowledging its limitations. It also discusses ethical and regulatory issues, including the lawyer's role in upholding the rule of law.

While labeled data is expensive to prepare, ever increasing amounts of unlabeled data is becoming widely available. In order to adapt to this phenomenon, several semi-supervised learning (SSL) algorithms, which learn from labeled as well as unlabeled data, have been developed. In a separate line of work, researchers have started to realize that graphs provide a natural way to represent data in a variety of domains. Graph-based SSL algorithms, which bring together these two lines of work, have been shown to outperform the state-of-the-art in many applications in speech processing, computer vision, natural language processing, and other areas of Artificial Intelligence. Recognizing this promising and emerging area of research, this synthesis lecture focuses on graph-based SSL algorithms (e.g., label propagation methods). Our hope is that after reading this book, the reader will walk away with the following: (1) an in-depth knowledge of the current state-of-the-art in graph-based SSL algorithms, and the ability to implement them; (2) the ability to decide on the suitability of graph-based SSL methods for a problem; and (3) familiarity with different applications where graph-based SSL methods have been successfully applied. Table of Contents: Introduction / Graph Construction / Learning and Inference / Scalability / Applications / Future Work / Bibliography / Authors' Biographies / Index

Exact Algorithms
Case-Based Reasoning

Advances in Case-Based Reasoning
Case-Based Reasoning on Images and Signals

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.

Case-based reasoning is a methodology with a long tradition in artificial intelligence that brings together reasoning and machine learning techniques to solve problems based on past experiences or cases. Given a problem to be solved, reasoning involves the use of methods to retrieve similar past cases in order to reuse their solution for the problem at hand. Once the problem has been solved, learning methods can be applied to improve the knowledge based on past experiences. In spite of being a broad methodology applied in industry and services, case-based reasoning has often been forgotten in both artificial intelligence and machine learning books. The aim of this book is to present a concise introduction to case-based reasoning providing the essential building blocks for the designing of case-based reasoning systems, as well as to bring together the main research lines in this field to encourage future students to solve current CBR challenges.

Judgment aggregation is a mathematical theory of collective decision-making. It concerns the methods whereby individual opinions about logically interconnected issues of interest can, or

cannot, be aggregated into one collective stance. Aggregation problems have traditionally been of interest for disciplines like economics and the political sciences, as well as philosophy, where judgment aggregation itself originates from, but have recently captured the attention of disciplines like computer science, artificial intelligence and multi-agent systems. Judgment aggregation has emerged in the last decade as a unifying paradigm for the formalization and understanding of aggregation problems. Still, no comprehensive presentation of the theory is available to date. This Synthesis Lecture aims at filling this gap presenting the key motivations, results, abstractions and techniques underpinning it. Table of Contents: Preface / Acknowledgments / Logic Meets Social Choice Theory / Basic Concepts / Impossibility / Coping with Impossibility / Manipulability / Aggregation Rules / Deliberation / Bibliography / Authors' Biographies / Index

Expert systems allow scientists to access, manage, and apply data and specialized knowledge from various disciplines to their own research. Expert Systems in Chemistry Research explains the general scientific basis and computational principles behind expert systems and demonstrates how they can improve the efficiency of scientific workflows and support decision-making processes. Focused initially on clarifying the fundamental concepts, limits, and drawbacks of using computer software to approach human decision making, the author also underscores the importance of putting theory into practice. The book highlights current capabilities for planning and monitoring experiments, scientific data management and interpretation, chemical characterization, problem solving, and methods for encoding chemical data. It also examines the challenges as well as requirements, strategies, and considerations for implementing expert systems effectively in an existing laboratory

software environment. Expert Systems in Chemistry Research covers various artificial intelligence technologies used to support expert systems, including nonlinear statistics, wavelet transforms, artificial neural networks, genetic algorithms, and fuzzy logic. This definitive text provides researchers, scientists, and engineers with a cornerstone resource for developing new applications in chemoinformatics, systems design, and other emerging fields.

From Foundations to Applications

Case-Based Reasoning Technology

Lifelong Machine Learning

Logic, Probability, and Computation

Statistical Relational Artificial Intelligence

Case-based ReasoningA Concise IntroductionMorgan & Claypool Publishers

This state-of-the-art survey presents a coherent summary of research and development in case-based reasoning (CBR) undertaken in Germany in recent years. The book opens with a general introduction to CBR presenting the basic ideas and concepts, setting the terminology, and looking at CBR from some new points of view. The main part of the book, consisting of nine chapters, is devoted to detailed presentations of CBR applications successfully performed in various areas. Among these application areas are decision and sales support, text processing, adaptation, planning, design, software engineering, tutoring systems, and medicine. The remaining

chapters present areas related to CBR as well as a glossary, a subject index and bibliography.

Is the Case-Based Reasoning scope complete and appropriately sized? How will you measure your Case-Based Reasoning effectiveness? How can the value of Case-Based Reasoning be defined? What is your formula for success in Case-Based Reasoning ? Will Case-Based Reasoning deliverables need to be tested and, if so, by whom? This powerful Case-Based Reasoning self-assessment will make you the entrusted Case-Based Reasoning domain adviser by revealing just what you need to know to be fluent and ready for any Case-Based Reasoning challenge. How do I reduce the effort in the Case-Based Reasoning work to be done to get problems solved? How can I ensure that plans of action include every Case-Based Reasoning task and that every Case-Based Reasoning outcome is in place? How will I save time investigating strategic and tactical options and ensuring Case-Based Reasoning costs are low? How can I deliver tailored Case-Based Reasoning advice instantly with structured going-forward plans? There's no better guide through these mind-expanding questions than acclaimed best-selling author Gerard Blokdyk. Blokdyk ensures all Case-Based Reasoning essentials are covered, from every angle: the Case-Based Reasoning self-assessment shows succinctly and clearly that what needs to be clarified to

organize the required activities and processes so that Case-Based Reasoning outcomes are achieved. Contains extensive criteria grounded in past and current successful projects and activities by experienced Case-Based Reasoning practitioners. Their mastery, combined with the easy elegance of the self-assessment, provides its superior value to you in knowing how to ensure the outcome of any efforts in Case-Based Reasoning are maximized with professional results. Your purchase includes access details to the Case-Based Reasoning self-assessment dashboard download which gives you your dynamically prioritized projects-ready tool and shows you exactly what to do next. Your exclusive instant access details can be found in your book. You will receive the following contents with New and Updated specific criteria: - The latest quick edition of the book in PDF - The latest complete edition of the book in PDF, which criteria correspond to the criteria in... - The Self-Assessment Excel Dashboard - Example pre-filled Self-Assessment Excel Dashboard to get familiar with results generation - In-depth and specific Case-Based Reasoning Checklists - Project management checklists and templates to assist with implementation INCLUDES LIFETIME SELF ASSESSMENT UPDATES Every self assessment comes with Lifetime Updates and Lifetime Free Updated Books. Lifetime Updates is an industry-first feature which allows you to receive verified self assessment updates,

ensuring you always have the most accurate information at your fingertips. This book constitutes the proceedings of the International Conference on Cognitive Computing, ICCC 2019, held as part of SCF 2019, in San Diego, CA, USA, in June 2019. The 14 full and 3 short papers presented in this volume were carefully reviewed and selected from 42 submissions. The papers cover all aspects of Sensing Intelligence (SIJ as a Service (SlaaS). Cognitive Computing is a sensing-driven computing (SDC) scheme that explores and integrates intelligence from all types of senses in various scenarios and solution contexts

Data Fusion Support to Activity-Based Intelligence

Adversarial Machine Learning

Graph-Based Semi-Supervised Learning

Strategic Voting

From Prediction to Action

Judgment aggregation is a mathematical theory of collective decision-making. It concerns the methods whereby individual opinions about logically interconnected issues of interest can, or cannot, be aggregated into one collective stance. Aggregation problems have traditionally been of interest for disciplines like economics and the political sciences, as well as philosophy, where judgment aggregation itself originates from, but have

recently captured the attention of disciplines like computer science, artificial intelligence and multi-agent systems. Judgment aggregation has emerged in the last decade as a unifying paradigm for the formalization and understanding of aggregation problems. Still, no comprehensive presentation of the theory is available to date. This Synthesis Lecture aims at filling this gap presenting the key motivations, results, abstractions and techniques underpinning it.

This book is the ?rst edited book that deals with the special topic of signals and images within case-based reasoning (CBR). Signal-interpreting systems are becoming increasingly popular in medical, industrial, ecological, biotechnological and many other applications. Existing statisticalandknowledge-basedtechniqueslackrobustness,accuracy,and?-ibility. New strategies are needed that can adapt to changing environmental conditions, signal variation, user needs and process requirements. Introducing CBRstrategiesintosignal-interpretingsystemscansatisfytheserequirements. CBR can be used to control the signal-processing process in all phases of a signal-interpreting system to derive information of the highest possible qu- ity. Beyond this CBR o?ers di?erent learning capabilities, for all phases of a signal-interpretingsystem,thatsatisfydi?erentneedsduringthedevelopment process

of a signal-interpreting system. In the outline of this book we summarize under the term "signal" signals of 1-dimensional, 2-dimensional or 3-dimensional nature. The unique data and the necessary computation techniques require ext- ordinary case representations, similarity measures and CBR strategies to be utilised.

Signalinterpretation(1D,2D,or3Dsignalinterpretation)istheprocessof mapping the numerical representation of a signal into logical representations suitable for signal descriptions. A signal-interpreting system must be able to extract symbolic features from the raw data e.g., the image (e.g., irregular structure inside the nodule, area of calci?cation, and sharp margin). This is a complex process; the signal passes through several general processing steps before the ?nal symbolic description is obtained. The structure of the book is divided into a theoretical part and into an application-oriented part.

Case-based reasoning is a methodology with a long tradition in artificial intelligence that brings together reasoning and machine learning techniques to solve problems based on past experiences or cases. Given a problem to be solved, reasoning involves the use of methods to retrieve similar past cases in order to reuse their solution for the problem at hand. Once the problem has been solved, learning methods can be applied to

improve the knowledge based on past experiences. In spite of being a broad methodology applied in industry and services, case-based reasoning has often been forgotten in both artificial intelligence and machine learning books. The aim of this book is to present a concise introduction to case-based reasoning providing the essential building blocks for the design of case-based reasoning systems, as well as to bring together the main research lines in this field to encourage students to solve current CBR challenges.

This book provides a tutorial introduction to modern techniques for representing and reasoning about qualitative preferences with respect to a set of alternatives. The syntax and semantics of several languages for representing preference languages, including CP-nets, TCP-nets, CI-nets, and CP-theories, are reviewed. Some key problems in reasoning about preferences are introduced, including determining whether one alternative is preferred to another, or whether they are equivalent, with respect to a given set of preferences. These tasks can be reduced to model checking in temporal logic. Specifically, an induced preference graph that represents a given set of preferences can be efficiently encoded using a Kripke Structure for Computational Tree Logic (CTL). One can translate preference queries with respect to a set of preferences into an equivalent set of formulae in

CTL, such that the CTL formula is satisfied whenever the preference query holds. This allows us to use a model checker to reason about preferences, i.e., answer preference queries, and to obtain a justification as to why a preference query is satisfied (or not) with respect to a set of preferences. This book defines the notions of the equivalence of two sets of preferences, including what it means for one set of preferences to subsume another, and shows how to answer preferential equivalence and subsumption queries using model checking. Furthermore, this book demontrates how to generate alternatives ordered by preference, along with providing ways to deal with inconsistent preference specifications. A description of CRISNER—an open source software implementation of the model checking approach to qualitative preference reasoning in CP-nets, TCP-nets, and CP-theories is included, as well as examples illustrating its use.

Introduction to Symbolic Plan and Goal Recognition

Case-Based Reasoning Research and Development

Theories, Methods, and Applications

5th International Conference on Case-Based Reasoning, ICCBR 2003, Trondheim, Norway, June 23-26, 2003, Proceedings

Reasoning with Probabilistic and Deterministic Graphical Models

The International Conference on Case-Based Reasoning (ICCBR) is the

pree- nent international meeting on case-based reasoning (CBR). ICCBR 2003 (http://www.iccbr.org/iccbr03/)isthe?fthinthisseriesofbiennialinter-tional conferences highlighting the most signi?cant contributions to the ?eld of CBR.TheconferencetookplacefromJune23throughJune26,2003attheN-wegian University of Science and Technology in Trondheim, Norway. Previous ICCBR conferences have been held in Vancouver, Canada (2001), Seeon, G- many (1999), Providence, Rhode Island, USA (1997), and Sesimbra, Portugal (1995). Day 1 of ICCBR 2003, Industry Day, provided hands-on experiences utilizing CBR in cutting-edge knowledge-management applications (e.g., help-desks,- business, and diagnostics). Day 2 featured topical workshops on CBR in the healthsciences,theimpactoflife-cyclemodelsonCBRsystems,mixed-initiative CBR, predicting time series with cases, and providing assistance with structured vs. unstructured cases. Days 3 and 4 comprised presentations and posters on theoretical and applied CBR research and deployed CBR applications, as well as invited talks from three distinguished scholars: David Leake, Indiana University, H ? ector Munoz-Avila, ? Lehigh University, and Ellen Rilo?, University of Utah. The presentations and posters covered

a wide range of CBR topics of in- rest both to practitioners and researchers, including case representation, si- larity, retrieval, adaptation, case library maintenance, multi-agent collaborative systems, data mining, soft computing, recommender systems, knowledge ma- gement, legal reasoning, software reuse and music.

The increasing abundance of large high-quality datasets, combined with significant technical advances over the last several decades have made machine learning into a major tool employed across a broad array of tasks including vision, language, finance, and security. However, success has been accompanied with important new challenges: many applications of machine learning are adversarial in nature. Some are adversarial because they are safety critical, such as autonomous driving. An adversary in these applications can be a malicious party aimed at causing congestion or accidents, or may even model unusual situations that expose vulnerabilities in the prediction engine. Other applications are adversarial because their task and/or the data they use are. For example, an important class of problems in security involves detection, such as malware, spam, and intrusion detection. The use of machine learning for detecting malicious entities creates an incentive among adversaries to evade

detection by changing their behavior or the content of malicious objects they develop. The field of adversarial machine learning has emerged to study vulnerabilities of machine learning approaches in adversarial settings and to develop techniques to make learning robust to adversarial manipulation. This book provides a technical overview of this field. After reviewing machine learning concepts and approaches, as well as common use cases of these in adversarial settings, we present a general categorization of attacks on machine learning. We then address two major categories of attacks and associated defenses: decision-time attacks, in which an adversary changes the nature of instances seen by a learned model at the time of prediction in order to cause errors, and poisoning or training time attacks, in which the actual training dataset is maliciously modified. In our final chapter devoted to technical content, we discuss recent techniques for attacks on deep learning, as well as approaches for improving robustness of deep neural networks. We conclude with a discussion of several important issues in the area of adversarial learning that in our view warrant further research. Given the increasing interest in the area of adversarial machine learning, we hope this book provides readers with the tools necessary to successfully engage in research and

practice of machine learning in adversarial settings.

From its inception, artificial intelligence (AI) has had a rather ambivalent relationship with humans—swinging between their augmentation and replacement. Now, as AI technologies enter our everyday lives at an ever-increasing pace, there is a greater need for AI systems to work synergistically with humans. One critical requirement for such synergistic human?AI interaction is that the AI systems' behavior be explainable to the humans in the loop. To do this effectively, AI agents need to go beyond planning with their own models of the world, and take into account the mental model of the human in the loop. At a minimum, AI agents need approximations of the human's task and goal models, as well as the human's model of the AI agent's task and goal models. The former will guide the agent to anticipate and manage the needs, desires and attention of the humans in the loop, and the latter allow it to act in ways that are interpretable to humans (by conforming to their mental models of it), and be ready to provide customized explanations when needed. The authors draw from several years of research in their lab to discuss how an AI agent can use these mental models to either conform to human expectations or change those expectations through explanatory communication. While the

focus of the book is on cooperative scenarios, it also covers how the same mental models can be used for obfuscation and deception. The book also describes several real-world application systems for collaborative decision-making that are based on the framework and techniques developed here. Although primarily driven by the authors' own research in these areas, every chapter will provide ample connections to relevant research from the wider literature. The technical topics covered in the book are self-contained and are accessible to readers with a basic background in AI.

An intelligent agent interacting with the real world will encounter individual people, courses, test results, drugs prescriptions, chairs, boxes, etc., and needs to reason about properties of these individuals and relations among them as well as cope with uncertainty. Uncertainty has been studied in probability theory and graphical models, and relations have been studied in logic, in particular in the predicate calculus and its extensions. This book examines the foundations of combining logic and probability into what are called relational probabilistic models. It introduces representations, inference, and learning techniques for probability, logic, and their combinations. The book focuses on two representations in detail: Markov logic networks, a relational extension of undirected graphical models and

weighted first-order predicate calculus formula, and Problog, a probabilistic extension of logic programs that can also be viewed as a Turing-complete relational extension of Bayesian networks.

Case-Based Reasoning a Clear and Concise Reference

Representing and Reasoning with Qualitative Preferences

Judgment Aggregation

Positive Unlabeled Learning

A Primer

Human decision-making often transcends our formal models of "rationality." Designing intelligent agents that interact proficiently with people necessitates the modeling of human behavior and the prediction of their decisions. In this book, we explore the task of automatically predicting human decision-making and its use in designing intelligent human-aware automated computer systems of varying natures—from purely conflicting interaction settings (e.g., security and games) to fully cooperative interaction settings (e.g., autonomous driving and personal robotic assistants). We explore the techniques, algorithms, and empirical methodologies for meeting the challenges that arise from the above tasks and illustrate major benefits from the use of these computational solutions in real-world application domains such as security, negotiations,

argumentative interactions, voting systems, autonomous driving, and games. The book presents both the traditional and classical methods as well as the most recent and cutting edge advances, providing the reader with a panorama of the challenges and solutions in predicting human decision-making. heterogeneous graphs. Further, the book introduces different applications of NE such as recommendation and information diffusion prediction. Finally, the book concludes the methods and applications and looks forward to the future directions.

The 2001 International Conference on Case-Based Reasoning (ICCBR 2001, www.iccbr.org/iccbr01), the fourth in the biennial ICCBR series (1995 in Sesimbra, Portugal; 1997 in Providence, Rhode Island (USA); 1999 in Seeon, Germany), was held during 30 July – 2 August 2001 in Vancouver, Canada. ICCBR is the premier international forum for researchers and practitioners of case based reasoning (CBR). The objectives of this meeting were to nurture significant, relevant advances made in this field (both in research and application), communicate them among all attendees, inspire future advances, and continue to support the vision that CBR is a valuable process in many research disciplines, both computational and otherwise. ICCBR 2001 was the first ICCBR meeting held on the Pacific coast, and we used the setting of

beautiful Vancouver as an opportunity to enhance participation from the Pacific Rim communities, which contributed 28% of the submissions. During this meeting, we were fortunate to host invited talks by Ralph Bergmann, Ken Forbus, Jaiwei Han, Ramon López de Mántaras, and Manuela Veloso. Their contributions ensured a stimulating meeting; we thank them all.

Many real-world decision problems have multiple objectives. For example, when choosing a medical treatment plan, we want to maximize the efficacy of the treatment, but also minimize the side effects. These objectives typically conflict, e.g., we can often increase the efficacy of the treatment, but at the cost of more severe side effects. In this book, we outline how to deal with multiple objectives in decision-theoretic planning and reinforcement learning algorithms. To illustrate this, we employ the popular problem classes of multi-objective Markov decision processes (MOMDPs) and multi-objective coordination graphs (MO-CoGs). First, we discuss different use cases for multi-objective decision making, and why they often necessitate explicitly multi-objective algorithms. We advocate a utility-based approach to multi-objective decision making, i.e., that what constitutes an optimal solution to a multi-objective decision problem should be derived from the available information about user utility. We show how different assumptions about user utility and

what types of policies are allowed lead to different solution concepts, which we outline in a taxonomy of multi-objective decision problems. Second, we show how to create new methods for multi-objective decision making using existing single-objective methods as a basis. Focusing on planning, we describe two ways to creating multi-objective algorithms: in the inner loop approach, the inner workings of a single-objective method are adapted to work with multi-objective solution concepts; in the outer loop approach, a wrapper is created around a single-objective method that solves the multi-objective problem as a series of single-objective problems. After discussing the creation of such methods for the planning setting, we discuss how these approaches apply to the learning setting. Next, we discuss three promising application domains for multi-objective decision making algorithms: energy, health, and infrastructure and transportation. Finally, we conclude by outlining important open problems and promising future directions.

8th European Conference, ECCBR 2006, Fethiye, Turkey, September 4-7, 2006, Proceedings

Artificial Intelligence and the Legal Profession

Tools and Applications

Multi-Objective Decision Making

Third International Conference, Held as Part of the Services Conference Federation, SCF 2019, San Diego, CA, USA, June 25–30, 2019, Proceedings

Graphs are useful data structures in complex real-life applications such as modeling physical systems, learning molecular fingerprints, controlling traffic networks, and recommending friends in social networks. However, these tasks require dealing with non-Euclidean graph data that contains rich relational information between elements and cannot be well handled by traditional deep learning models (e.g., convolutional neural networks (CNNs) or recurrent neural networks (RNNs)). Nodes in graphs usually contain useful feature information that cannot be well addressed in most unsupervised representation learning methods (e.g., network embedding methods). Graph neural networks (GNNs) are proposed to combine the feature information and the graph structure to learn better representations on graphs via feature propagation and aggregation. Due to its convincing performance and high interpretability, GNN has recently become a widely applied graph analysis tool. This book provides a comprehensive introduction to the basic concepts, models, and applications of graph neural networks. It starts with the introduction of the vanilla GNN model. Then several variants of the vanilla model are introduced such as graph convolutional networks, graph recurrent networks, graph attention networks, graph residual networks, and several general frameworks. Variants for different graph types and advanced training methods are also included. As for the applications of GNNs, the book categorizes them

into structural, non-structural, and other scenarios, and then it introduces several typical models on solving these tasks. Finally, the closing chapters provide GNN open resources and the outlook of several future directions.

This is a technical overview of the field of adversarial machine learning which has emerged to study vulnerabilities of machine learning approaches in adversarial settings and to develop techniques to make learning robust to adversarial manipulation. After reviewing machine learning concepts and approaches, as well as common use cases of these in adversarial settings, we present a general categorization of attacks on machine learning. We then address two major categories of attacks and associated defenses: decision-time attacks, in which an adversary changes the nature of instances seen by a learned model at the time of prediction in order to cause errors, and poisoning or training time attacks, in which the actual training dataset is maliciously modified. In our final chapter devoted to technical content, we discuss recent techniques for attacks on deep learning, as well as approaches for improving robustness of deep neural networks. We conclude with a discussion of several important issues in the area of adversarial learning that in our view warrant further research. The increasing abundance of large high-quality datasets, combined with significant technical advances over the last several decades have made machine learning into a major tool employed across a broad array of tasks including vision, language, finance, and security. However, success has been accompanied with important new challenges: many

applications of machine learning are adversarial in nature. Some are adversarial because they are safety critical, such as autonomous driving. An adversary in these applications can be a malicious party aimed at causing congestion or accidents, or may even model unusual situations that expose vulnerabilities in the prediction engine. Other applications are adversarial because their task and/or the data they use are. For example, an important class of problems in security involves detection, such as malware, spam, and intrusion detection. The use of machine learning for detecting malicious entities creates an incentive among adversaries to evade detection by changing their behavior or the content of malicious objects they develop. Given the increasing interest in the area of adversarial machine learning, we hope this book provides readers with the tools necessary to successfully engage in research and practice of machine learning in adversarial settings.

Graphical models (e.g., Bayesian and constraint networks, influence diagrams, and Markov decision processes) have become a central paradigm for knowledge representation and reasoning in both artificial intelligence and computer science in general. These models are used to perform many reasoning tasks, such as scheduling, planning and learning, diagnosis and prediction, design, hardware and software verification, and bioinformatics. These problems can be stated as the formal tasks of constraint satisfaction and satisfiability, combinatorial optimization, and probabilistic inference. It is well known that the tasks are computationally hard, but research during

the past three decades has yielded a variety of principles and techniques that significantly advanced the state of the art. In this book we provide comprehensive coverage of the primary exact algorithms for reasoning with such models. The main feature exploited by the algorithms is the model's graph. We present inference-based, message-passing schemes (e.g., variable-elimination) and search-based, conditioning schemes (e.g., cycle-cutset conditioning and AND/OR search). Each class possesses distinguished characteristics and in particular has different time vs. space behavior. We emphasize the dependence of both schemes on few graph parameters such as the treewidth, cycle-cutset, and (the pseudo-tree) height. We believe the principles outlined here would serve well in moving forward to approximation and anytime-based schemes. The target audience of this book is researchers and students in the artificial intelligence and machine learning area, and beyond.

Intelligent systems often depend on data provided by information agents, for example, sensor data or crowdsourced human computation. Providing accurate and relevant data requires costly effort that agents may not always be willing to provide. Thus, it becomes important not only to verify the correctness of data, but also to provide incentives so that agents that provide high-quality data are rewarded while those that do not are discouraged by low rewards. We cover different settings and the assumptions they admit, including sensing, human computation, peer grading, reviews, and predictions. We survey different incentive mechanisms, including proper scoring rules, prediction

markets and peer prediction, Bayesian Truth Serum, Peer Truth Serum, Correlated Agreement, and the settings where each of them would be suitable. As an alternative, we also consider reputation mechanisms. We complement the game-theoretic analysis with practical examples of applications in prediction platforms, community sensing, and peer grading.

Expert Systems in Chemistry Research

A Planning Perspective

Cognitive Computing – ICCC 2019

Introduction to Graph Neural Networks

Transfer Learning for Multiagent Reinforcement Learning Systems

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.
This book constitutes the refereed proceedings of the 8th European Conference on Case-Based Reasoning, ECCBR 2004, held in Fethiye, Turkey in September 2006. The book presents 31 revised full papers and 5 revised application papers together with 2 invited papers and 2 abstracts of invited talks. The coverage represents snapshot of current current issues in case-based reasoning, ranging from theoretical and

methodological issues to advanced applications in various fields.

Lifelong Machine Learning, Second Edition

Graph Representation Learning

Explainable Human-AI Interaction

Case-based Reasoning

Second Edition