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

Second Edition
Eliciting Truthful Information
Computer Science And Technology - Proceedings Of The International Conference (Cst2016)
Case-Based Reasoning
A Concise Introduction to Models and Methods for Automated Planning
Topics in Case-based Reasoning
Multi-Objective Decision Making
Transfer Learning for Multiagent Reinforcement Learning Systems
Judgment Aggregation
Strategic Voting
A Textbook
Data Fusion Support to Activity-Based Intelligence
Reasoning with Complex Cases
Metric Learning
Graph Representation Learning
General Game Playing
Introduction to Graph Neural Networks
Third International Conference, Held as Part of the Services Conference Federation, SCF 2019, San Diego, CA, USA, June 25–30, 2019, Proceedings
Artificial Intelligence and the Legal Profession
Case-Based Reasoning
Case-Based Reasoning a Clear and Concise Reference
An Introduction to Constraint-Based Temporal Reasoning
A Planning Perspective
Federated Learning
Reasoning with Probabilistic and Deterministic Graphical Models
Introduction to Symbolic Plan and Goal Recognition
Cognitive Computing – ICCC 2019
Introduction to Intelligent Systems in Traffic and Transportation
Case-based Reasoning
Exact Algorithms
Case-Based Reasoning
Explainable Human-AI Interaction
18th International Conference, ICCBR 2010, Alessandria, Italy, July 19-22, 2010

Proceedings
Case-Based Reasoning Research and Development
Statistical Relational Artificial Intelligence
Analogies and Theories
Formal Models of Reasoning
From Prediction to Action

EATON AGUILAR

Second Edition
Bloomsbury Publishing
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.
Eliciting Truthful Information Morgan & Claypool Publishers
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
Computer Science And Technology - Proceedings Of The International Conference (Cst2016)
World Scientific
Urban mobility is not only one of the pillars of modern economic systems, but also a key issue in the quest for equality of opportunity, once it can improve access to other services. Currently, however, there are a number of negative issues related to traffic, especially in mega-cities, such as economical issues (cost of opportunity caused by delays), environmental (externalities related to emissions of pollutants), and social (traffic accidents). Solutions to these issues are more and

more closely tied to information and communication technology. Indeed, a search in the technical literature (using the keyword ``urban traffic" to filter out articles on data network traffic) retrieved the following number of articles (as of December 3, 2013): 9,443 (ACM Digital Library), 26,054 (Scopus), and 1,730,000 (Google Scholar). Moreover, articles listed in the ACM query relate to conferences as diverse as MobiCom, CHI, PADS, and AAMAS. This means that there is a big and diverse community of computer scientists and computer engineers who tackle research that is connected to the development of intelligent traffic and transportation systems. It is also possible to see that this community is growing, and that research projects are getting more and more interdisciplinary. To foster the cooperation among the involved communities, this book aims at giving a broad introduction into the basic but relevant concepts related to transportation systems, targeting researchers and practitioners from computer science and information technology. In addition, the second part of the book gives a panorama of some of the most exciting and newest technologies, originating in computer science and computer engineering, that are now being employed in projects related to car-to-car communication, interconnected vehicles, car navigation, platooning, crowd sensing and sensor networks, among others. This material will also be of interest to engineers and researchers from the traffic and transportation community.

Case-Based Reasoning

Morgan & Claypool Publishers

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.

A Concise Introduction to Models and Methods for Automated Planning

Morgan & Claypool Publishers

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.

Topics in Case-based Reasoning Morgan & Claypool Publishers

Solving challenging computational problems involving time has been a critical component in the development of artificial intelligence systems almost since the inception of the field. This book provides a concise introduction to the core computational elements of temporal reasoning for use in AI systems for planning and scheduling, as well as systems that extract temporal information from data. It presents a survey of temporal frameworks based on constraints, both qualitative and quantitative, as well as of major temporal consistency techniques. The book also introduces the reader to more recent extensions to the core model that allow AI systems to explicitly represent temporal preferences and temporal uncertainty. This book is intended for students and researchers interested in constraint-based temporal reasoning. It provides a self-contained guide to the different representations of time, as well as examples of recent applications of time in AI systems.

Multi-Objective Decision Making Morgan & Claypool Publishers

Planning is the branch of Artificial Intelligence (AI) that seeks to automate reasoning about plans, most importantly the reasoning that goes into formulating a plan to achieve a given goal in a given situation. AI planning is model-based: a planning system takes as input a description (or model) of the initial situation, the actions available to change it, and the goal condition to output a plan composed of those actions that will accomplish the goal when executed from the initial situation. The Planning Domain Definition Language (PDDL) is a formal knowledge representation language designed to express planning models. Developed by the planning research community as a means of facilitating systems comparison, it has become a de-facto standard input language of many planning systems, although it is not the only modelling language for planning. Several variants of PDDL have emerged that capture planning problems of different natures and complexities, with a focus on deterministic problems. The purpose of this book is two-fold. First, we present a unified and current account of PDDL, covering the subsets of PDDL that express discrete, numeric,

temporal, and hybrid planning. Second, we want to introduce readers to the art of modelling planning problems in this language, through educational examples that demonstrate how PDDL is used to model realistic planning problems. The book is intended for advanced students and researchers in AI who want to dive into the mechanics of AI planning, as well as those who want to be able to use AI planning systems without an in-depth explanation of the algorithms and implementation techniques they use.

Transfer Learning for Multiagent Reinforcement Learning Systems Morgan & Claypool Publishers

Reasoning with Complex Cases emphasizes case retrieval methods based on structured cases as they are relevant for planning, configuration, and design, and provides a systematic view of the case reuse phase, centering on complex situations. So far, books on case-based reasoning considered comparatively simple situations only. This book is a coherent work, not a selection of separate contributions, and consists largely of original research results using examples taken from industrial design, biology, medicine, jurisprudence and other areas. Reasoning with Complex Cases is suitable as a secondary text for graduate-level courses on case-based reasoning and as a reference for practitioners applying conventional CBR systems or techniques.

Judgment Aggregation Morgan & Claypool Publishers

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. Strategic Voting Springer Science & Business Media Many machine learning algorithms require real-valued feature vectors of data instances as inputs. By projecting data into vector spaces, representation learning techniques have achieved promising performance in many areas such as computer vision and natural language processing. There is also a need to learn representations for discrete relational data, namely networks or graphs. Network Embedding (NE) aims at learning vector representations for each node or vertex in a network to encode the topologic structure. Due to its convincing performance and efficiency, NE has been widely applied in many network applications such as node classification and link prediction. This book provides a comprehensive introduction to the basic concepts, models, and applications of network representation learning (NRL). The book starts with an introduction to the background and rising of network embeddings as a general overview for readers. Then it introduces the development of NE techniques by presenting several representative methods on general graphs, as well as a unified NE framework based on matrix factorization. Afterward, it presents the variants of NE with additional information: NE for graphs with node attributes/contents/labels; and the variants with different characteristics: NE for community-structured/large-scale/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. Case-based ReasoningA Concise Introduction 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.

A Textbook Morgan & Claypool Publishers

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.
Data Fusion Support to Activity-Based Intelligence OUP Oxford
Case-based ReasoningA Concise IntroductionMorgan & Claypool Publishers
Reasoning with Complex Cases Morgan & Claypool Publishers
Planning is the model-based approach to autonomous behavior where the agent behavior is derived automatically from a model of the actions, sensors, and goals. The main challenges in planning are computational as all models, whether featuring uncertainty and feedback or not, are intractable in the worst case when represented in compact form. In this book, we look at a variety of models used in AI planning, and at the methods that have been developed for solving them. The goal is to provide a modern and coherent view of planning that is precise, concise, and mostly self-contained, without being shallow. For this, we make no attempt at covering the whole variety of planning approaches, ideas, and applications, and focus on the essentials. The target audience of the book are students and researchers interested in autonomous behavior and planning from an AI, engineering, or cognitive science perspective. Table of Contents: Preface / Planning and Autonomous Behavior / Classical Planning: Full Information and Deterministic Actions / Classical Planning: Variations and Extensions / Beyond Classical Planning: Transformations / Planning with Sensing: Logical Models / MDP Planning: Stochastic Actions and Full Feedback / POMDP Planning: Stochastic Actions and Partial Feedback / Discussion / Bibliography / Author's Biography
Metric Learning Springer
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.

Graph Representation Learning CRC Press

Similarity between objects plays an important role in both human cognitive processes and artificial systems for recognition and categorization. How to appropriately measure such similarities for a given task is crucial to the performance of many machine learning, pattern recognition and data mining methods. This book is devoted to metric learning, a set of techniques to automatically learn similarity and distance functions from data that has attracted a lot of interest in machine learning and related fields in the past ten years. In this book, we provide a thorough review of the metric learning literature that covers algorithms, theory and applications for both numerical and structured data. We first introduce relevant definitions and classic metric functions, as well as examples of their use in machine learning and data mining. We then review a wide range of metric learning algorithms, starting with the simple setting of linear distance and similarity learning. We show how one may scale-up these methods to very large amounts of training data. To go beyond the linear case, we discuss methods that learn nonlinear metrics or multiple linear metrics throughout the feature space, and review methods for more complex settings such as multi-task and semi-supervised learning. Although most of the existing work has focused on numerical data, we cover the literature on metric learning for structured data like strings, trees, graphs and time series. In the more technical part of the book, we present some recent statistical frameworks for analyzing the generalization performance in metric learning and derive results for some of the algorithms presented earlier. Finally, we illustrate the relevance of metric learning in real-world problems through a series of successful applications to computer vision, bioinformatics and information retrieval.

Artech House

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.

General Game Playing

Morgan & Claypool Publishers

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.

Introduction to Graph Neural Networks

Springer

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.

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

Morgan & Claypool Publishers

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.

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

• Sp2 Answers Key : click here