ESSENTIALS OF GAME THEORY A CONCISE MULTIDISCIPLINARY INTRODUCTION YOAV SHOHAM

Essentials of Game Theory

Game theory is the mathematical study of interaction among independent, self-interested agents. The audience for game theory has grown dramatically in recent years, and now spans disciplines as diverse as political science, biology, psychology, economics, linguistics, sociology, and computer science, among others. What has been missing is a relatively short introduction to the field covering the common basis that anyone with a professional interest in game theory is likely to require. Such a text would minimize notation, ruthlessly focus on essentials, and yet not sacrifice rigor. This Synthesis Lecture aims to fill this gap by providing a concise and accessible introduction to the field. It covers the main classes of games, their representations, and the main concepts used to analyze them.

Multiagent Systems

Multiagent systems combine multiple autonomous entities, each having diverging interests or different information. This overview of the field offers a computer science perspective, but also draws on ideas from game theory, economics, operations research, logic, philosophy and linguistics. It will serve as a reference for researchers in each of these fields, and be used as a text for advanced undergraduate or graduate courses. The authors emphasize foundations to create a broad and rigorous treatment of their subject, with thorough presentations of distributed problem solving, game theory, multiagent communication and learning, social choice, mechanism design, auctions, cooperative game theory, and modal logics of knowledge and belief. For each topic, basic concepts are introduced, examples are given, proofs of key results are offered, and algorithmic considerations are examined. An appendix covers background material in probability theory, classical logic, Markov decision processes and mathematical programming.

Essentials of Game Theory

Game theory is the mathematical study of interaction among independent, self-interested agents. The audience for game theory has grown dramatically in recent years, and now spans disciplines as diverse as political science, biology, psychology, economics, linguistics, sociology, and computer science, among others. What has been missing is a relatively short introduction to the field covering the common basis that anyone with a professional interest in game theory is likely to require. Such a text would minimize notation, ruthlessly focus on essentials, and yet not sacrifice rigor. This Synthesis Lecture aims to fill this gap by providing a concise and accessible introduction to the field. It covers the main classes of games, their representations, and the main concepts used to analyze them.

Game Theory and Machine Learning for Cyber Security

GAME THEORY AND MACHINE LEARNING FOR CYBER SECURITY Move beyond the foundations of machine learning and game theory in cyber security to the latest research in this cutting-edge field In Game Theory and Machine Learning for Cyber Security, a team of expert security researchers delivers a collection of central research contributions from both machine learning and game theory applicable to cybersecurity. The distinguished editors have included resources that address open research questions in game theory and machine learning applied to cyber security systems and examine the strengths and limitations of current game theoretic models for cyber security. Readers will explore the vulnerabilities of traditional machine learning algorithms and how they can be mitigated in an adversarial machine learning approach. The book offers a comprehensive suite of solutions to a broad range of technical issues in applying game theory and machine learning to solve cyber security challenges. Beginning with an introduction to foundational concepts in game theory, machine learning, cyber security, and cyber deception, the editors provide readers with resources that discuss the latest in hypergames, behavioral game theory, adversarial machine learning, generative adversarial networks, and multi-agent reinforcement learning. Readers will also enjoy: A thorough introduction to game theory for cyber deception, including scalable algorithms for identifying stealthy attackers in a game theoretic framework, honeypot allocation over attack graphs, and behavioral games for cyber deception An exploration of game theory for cyber security, including actionable game-theoretic adversarial intervention detection against advanced persistent threats Practical discussions of adversarial machine learning for cyber security, including adversarial machine learning in 5G security and machine learning-driven fault injection in cyber-physical systems In-depth examinations of generative models for cyber security Perfect for researchers, students, and experts in the fields of computer science and engineering, Game Theory and Machine Learning for Cyber Security is also an indispensable resource for industry professionals, military personnel, researchers, faculty, and students with an interest in cyber security.

Multiagent Systems, second edition

The new edition of an introduction to multiagent systems that captures the state of the art in both theory and practice, suitable as textbook or reference. Multiagent systems are made up of multiple interacting intelligent agents—computational entities to some degree autonomous and able to cooperate, compete, communicate, act flexibly, and exercise control over their behavior within the frame of their objectives. They are the enabling technology for a wide range of advanced applications relying on distributed and parallel processing of data, information, and knowledge relevant in domains ranging from industrial manufacturing to e-commerce to health care. This book offers a state-of-the-art introduction to multiagent systems, covering the field in both breadth and depth, and treating both theory and practice. It is suitable for classroom use or independent study. This second edition has been completely revised, capturing the tremendous developments in multiagent systems since the first edition appeared in 1999. Sixteen of the book's seventeen chapters were written for this edition; all chapters are by leaders in the field, with each author contributing to the broad base of knowledge and experience on which the book rests. The book covers basic concepts of computational agency from the perspective of both individual agents and agent organizations; communication among agents; coordination among agents; distributed cognition; development and engineering of multiagent systems; and background knowledge in logics and game theory. Each chapter includes references, many illustrations and examples, and exercises of varying degrees of difficulty. The chapters and the overall book are designed to be self-contained and understandable without additional material. Supplemental resources are available on the book's Web site. Contributors Rafael Bordini, Felix Brandt, Amit Chopra, Vincent Conitzer, Virginia Dignum, Jürgen Dix, Ed Durfee, Edith Elkind, Ulle Endriss, Alessandro Farinelli, Shaheen Fatima, Michael Fisher, Nicholas R. Jennings, Kevin Leyton-Brown, Evangelos Markakis, Lin Padgham, Julian Padget, Iyad Rahwan, Talal Rahwan, Alex Rogers, Jordi Sabater-Mir, Yoav Shoham, Munindar P. Singh, Kagan Tumer, Karl Tuyls, Wiebe van der Hoek, Laurent Vercouter, Meritxell Vinyals, Michael Winikoff, Michael Wooldridge, Shlomo Zilberstein

Computational Aspects of Cooperative Game Theory

Cooperative game theory is a branch of (micro-)economics that studies the behavior of self-interested agents in strategic settings where binding agreements among agents are possible. Our aim in this book is to present a survey of work on the computational aspects of cooperative game theory. We begin by formally defining transferable utility games in characteristic function form, and introducing key solution concepts such as the core and the Shapley value. We then discuss two major issues that arise when considering such games from a computational perspective: identifying compact representations for games, and the closely related problem of efficiently computing solution concepts for games. We survey several formalisms for cooperative games that have been proposed in the literature, including, for example, cooperative games defined on networks, as well as general compact representation schemes such as MC-nets and skill games. As a detailed case study, we consider weighted voting games: a widely-used and practically important class of cooperative games that inherently have a natural compact representation. We investigate the complexity of solution concepts for such games, and generalizations of them. We briefly discuss games with non-transferable utility and partition function games. We then overview algorithms for identifying welfare-maximizing coalition structures and methods used by rational agents to form coalitions (even under uncertainty), including bargaining algorithms. We conclude by considering some developing topics, applications, and future research directions.

Game Theory for Data Science

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.

Game Theory

This new edition is unparalleled in breadth of coverage, thoroughness of technical explanations and number of worked examples.

Playing with Scripture

This book puts a creative new reading of Hans-Georg Gadamer’s philosophical hermeneutics and literary genre theory to work on the problem of Scripture. Reading texts as Scripture brings two hermeneutical assumptions into tension: that the text will continually say something new and relevant to the present situation, and that the text has stability and authority over readers. Given how contested the Bible’s meaning is, how is it possible to ‘read Scripture’ as authoritative and relevant? Rather than anchor meaning in author, text or reader, Gadamer’s phenomenological model of hermeneutical experience as Spiel (‘play’) offers a dynamic, intersubjective account of how understanding happens, avoiding the dead end of the subjective–objective dichotomy. Modern genre theory addresses some of the criticisms of Gadamer, accounting for the different roles played by readers in different genres using the new term Lesespiel (‘reading game’). This is tested in three case studies of contested texts: the recontextualization of psalms in the book of Acts, the use of Hagar’s story (Genesis 16) in nineteenth-century debates over slavery and the troubling reception history of the rape and murder in Gibeah (Judges 19). In each study, the application of ancient text to contemporary situation is neither arbitrary, nor slavishly bound to tradition, but playful.

Lifelong Machine Learning

Lifelong Machine Learning (or Lifelong Learning) is an advanced machine learning paradigm that learns continuously, accumulates the knowledge learned in previous tasks, and uses it to help future learning. In the process, the learner becomes more and more knowledgeable and effective at learning. This learning ability is one of the hallmarks of human intelligence. However, 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. It makes no attempt to retain the learned knowledge and use it in future learning. Although this isolated learning paradigm has been very successful, it requires a large number of training examples, and is only suitable for well-defined and narrow tasks. In comparison, we humans can learn effectively with a few examples because we have accumulated so much knowledge in the past which enables us to learn with little data or effort. Lifelong learning aims to achieve this capability. As statistical machine learning matures, it is time to make a major effort to break the isolated learning tradition and to study lifelong learning to bring machine learning to new heights. Applications such as intelligent assistants, chatbots, and physical robots that interact with humans and systems in real-life environments are also calling for such lifelong learning capabilities. Without the ability to accumulate the learned knowledge and use it to learn more knowledge incrementally, a system will probably never be truly intelligent. This book serves as an introductory text and survey to lifelong learning.

Lifelong Machine Learning, Second Edition

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.

Federated Learning

How is it possible to allow multiple data owners to collaboratively train and use a shared prediction model while keeping all the local training data private? Traditional machine learning approaches need to combine all data at one location, typically a data center, which may very well violate the laws on user privacy and data confidentiality. Today, many parts of the world demand that technology companies treat user data carefully according to user-privacy laws. The European Union's General Data Protection Regulation (GDPR) is a prime example. In this book, we describe how federated machine learning addresses this problem with novel solutions combining distributed machine learning, cryptography and security, and incentive mechanism design based on economic principles and game theory. We explain different types of privacy-preserving machine learning solutions and their technological backgrounds, and highlight some representative practical use cases. We show how federated learning can become the foundation of next-generation machine learning that caters to technological and societal needs for responsible AI development and application.

Active Learning

The key idea behind active learning is that a machine learning algorithm can perform better with less training if it is allowed to choose the data from which it learns. An active learner may pose \"queries,\" usually in the form of unlabeled data instances to be labeled by an \"oracle\" (e.g., a human annotator) that already understands the nature of the problem. This sort of approach is well-motivated in many modern machine learning and data mining applications, where unlabeled data may be abundant or easy to come by, but training labels are difficult, time-consuming, or expensive to obtain. This book is a general introduction to active learning. It outlines several scenarios in which queries might be formulated, and details many query selection algorithms which have been organized into four broad categories, or \"query selection frameworks.\" We also touch on some of the theoretical foundations of active learning, and conclude with an overview of the strengths and weaknesses of these approaches in practice, including a summary of ongoing work to address these open challenges and opportunities. Table of Contents: Automating Inquiry / Uncertainty Sampling / Searching Through the Hypothesis Space / Minimizing Expected Error and Variance / Exploiting Structure in Data / Theory / Practical Considerations

Representations and Techniques for 3D Object Recognition and Scene Interpretation

One of the grand challenges of artificial intelligence is to enable computers to interpret 3D scenes and objects from imagery. This book organizes and introduces major concepts in 3D scene and object representation and inference from still images, with a focus on recent efforts to fuse models of geometry and perspective with statistical machine learning. The book is organized into three sections: (1) Interpretation of Physical Space; (2) Recognition of 3D Objects; and (3) Integrated 3D Scene Interpretation. The first discusses representations of spatial layout and techniques to interpret physical scenes from images. The second section introduces representations for 3D object categories that account for the intrinsically 3D nature of objects and provide robustness to change in viewpoints. The third section discusses strategies to unite inference of scene geometry and object pose and identity into a coherent scene interpretation. Each section broadly surveys important ideas from cognitive science and artificial intelligence research, organizes and discusses key concepts and techniques from recent work in computer vision, and describes a few sample approaches in detail. Newcomers to computer vision will benefit from introductions to basic concepts, such as single-view geometry and image classification, while experts and novices alike may find inspiration from the book's organization and discussion of the most recent ideas in 3D scene understanding and 3D object recognition. Specific topics include: mathematics of perspective geometry; visual elements of the physical scene, structural 3D scene representations; techniques and features for image and region categorization; historical perspective, computational models, and datasets and machine learning techniques for 3D object recognition; inferences of geometrical attributes of objects, such as size and pose; and probabilistic and feature-passing approaches for contextual reasoning about 3D objects and scenes. Table of Contents: Background on 3D Scene Models / Single-view Geometry / Modeling the Physical Scene / Categorizing Images and Regions / Examples of 3D Scene Interpretation / Background on 3D Recognition / Modeling 3D Objects / Recognizing and Understanding 3D Objects / Examples of 2D 1/2 Layout Models / Reasoning about Objects and Scenes / Cascades of Classifiers / Conclusion and Future Directions

Judgment Aggregation

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.

Human Computation

Human computation is a new and evolving research area that centers around harnessing human intelligence to solve computational problems that are beyond the scope of existing Artificial Intelligence (AI) algorithms. With the growth of the Web, human computation systems can now leverage the abilities of an unprecedented number of people via the Web to perform complex computation. There are various genres of human computation applications that exist today. Games with a purpose (e.g., the ESP Game) specifically target online gamers who generate useful data (e.g., image tags) while playing an enjoyable game. Crowdsourcing marketplaces (e.g., Amazon Mechanical Turk) are human computation systems that coordinate workers to perform tasks in exchange for monetary rewards. In identity verification tasks, users perform computation in order to gain access to some online content; an example is reCAPTCHA, which leverages millions of users who solve CAPTCHAs every day to correct words in books that optical character recognition (OCR) programs fail to recognize with certainty. This book is aimed at achieving four goals: (1) defining human computation as a research area; (2) providing a comprehensive review of existing work; (3) drawing connections to a wide variety of disciplines, including AI, Machine Learning, HCI, Mechanism/Market Design and Psychology, and capturing their unique perspectives on the core research questions in human computation; and (4) suggesting promising research directions for the future. Table of Contents: Introduction / Human Computation Algorithms / Aggregating Outputs / Task Routing / Understanding Workers and Requesters / The Art of Asking Questions / The Future of Human Computation

Planning with Markov Decision Processes

Markov Decision Processes (MDPs) are widely popular in Artificial Intelligence for modeling sequential decision-making scenarios with probabilistic dynamics. They are the framework of choice when designing an intelligent agent that needs to act for long periods of time in an environment where its actions could have uncertain outcomes. MDPs are actively researched in two related subareas of AI, probabilistic planning and reinforcement learning. Probabilistic planning assumes known models for the agent's goals and domain dynamics, and focuses on determining how the agent should behave to achieve its objectives. On the other hand, reinforcement learning additionally learns these models based on the feedback the agent gets from the environment. This book provides a concise introduction to the use of MDPs for solving probabilistic planning problems, with an emphasis on the algorithmic perspective. It covers the whole spectrum of the field, from the basics to state-of-the-art optimal and approximation algorithms. We first describe the theoretical foundations of MDPs and the fundamental solution techniques for them. We then discuss modern optimal algorithms based on heuristic search and the use of structured representations. A major focus of the book is on the numerous approximation schemes for MDPs that have been developed in the AI literature. These include determinization-based approaches, sampling techniques, heuristic functions, dimensionality reduction, and hierarchical representations. Finally, we briefly introduce several extensions of the standard MDP classes that model and solve even more complex planning problems. Table of Contents: Introduction / MDPs / Fundamental Algorithms / Heuristic Search Algorithms / Symbolic Algorithms / Approximation Algorithms / Advanced Notes

Data Integration

Data integration is a critical problem in our increasingly interconnected but inevitably heterogeneous world. There are numerous data sources available in organizational databases and on public information systems like the World Wide Web. Not surprisingly, the sources often use different vocabularies and different data structures, being created, as they are, by different people, at different times, for different purposes. The goal of data integration is to provide programmatic and human users with integrated access to multiple, heterogeneous data sources, giving each user the illusion of a single, homogeneous database designed for his or her specific need. The good news is that, in many cases, the data integration process can be automated. This book is an introduction to the problem of data integration and a rigorous account of one of the leading approaches to solving this problem, viz., the relational logic approach. Relational logic provides a theoretical framework for discussing data integration. Moreover, in many important cases, it provides algorithms for solving the problem in a computationally practical way. In many respects, relational logic does for data integration what relational algebra did for database theory several decades ago. A companion web site provides interactive demonstrations of the algorithms. Table of Contents: Preface / Interactive Edition / Introduction / Basic Concepts / Query Folding / Query Planning / Master Schema Management / Appendix / References / Index / Author Biography Don't have access? Recommend our Synthesis Digital Library to your library or purchase a personal subscription. Email info@morganclaypool.com for details.

Algorithms for Reinforcement Learning

Reinforcement learning is a learning paradigm concerned with learning to control a system so as to maximize a numerical performance measure that expresses a long-term objective. What distinguishes reinforcement learning from supervised learning is that only partial feedback is given to the learner about the learner's predictions. Further, the predictions may have long term effects through influencing the future state of the controlled system. Thus, time plays a special role. The goal in reinforcement learning is to develop efficient learning algorithms, as well as to understand the algorithms' merits and limitations. Reinforcement learning is of great interest because of the large number of practical applications that it can be used to address, ranging from problems in artificial intelligence to operations research or control engineering. In this book, we focus on those algorithms of reinforcement learning that build on the powerful theory of dynamic programming. We give a fairly comprehensive catalog of learning problems, describe the core ideas, note a large number of state of the art algorithms, followed by the discussion of their theoretical properties and limitations. Table of Contents: Markov Decision Processes / Value Prediction Problems / Control / For Further Exploration

Applying Reinforcement Learning on Real-World Data with Practical Examples in Python

Reinforcement learning is a powerful tool in artificial intelligence in which virtual or physical agents learn to optimize their decision making to achieve long-term goals. In some cases, this machine learning approach can save programmers time, outperform existing controllers, reach super-human performance, and continually adapt to changing conditions. This book argues that these successes show reinforcement learning can be adopted successfully in many different situations, including robot control, stock trading, supply chain optimization, and plant control. However, reinforcement learning has traditionally been limited to applications in virtual environments or simulations in which the setup is already provided. Furthermore, experimentation may be completed for an almost limitless number of attempts risk-free. In many real-life tasks, applying reinforcement learning is not as simple as (1) data is not in the correct form for reinforcement learning, (2) data is scarce, and (3) automation has limitations in the real-world. Therefore, this book is written to help academics, domain specialists, and data enthusiast alike to understand the basic principles of applying reinforcement learning to real-world problems. This is achieved by focusing on the process of taking practical examples and modeling standard data into the correct form required to then apply basic agents. To further assist with readers gaining a deep and grounded understanding of the approaches, the book shows hand-calculated examples in full and then how this can be achieved in a more automated manner with code. For decision makers who are interested in reinforcement learning as a solution but are not technically proficient we include simple, non-technical examples in the introduction and case studies section. These provide context of what reinforcement learning offer but also the challenges and risks associated with applying it in practice. Specifically, the book illustrates the differences between reinforcement learning and other machine learning approaches as well as how well-known companies have found success using the approach to their problems.

Strategic Voting

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 becomes 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.

Markov Logic

Most subfields of computer science have an interface layer via which applications communicate with the infrastructure, and this is key to their success (e.g., the Internet in networking, the relational model in databases, etc.). So far this interface layer has been missing in AI. First-order logic and probabilistic graphical models each have some of the necessary features, but a viable interface layer requires combining both. Markov logic is a powerful new language that accomplishes this by attaching weights to first-order formulas and treating them as templates for features of Markov random fields. Most statistical models in wide use are special cases of Markov logic, and first-order logic is its infinite-weight limit. Inference algorithms for Markov logic combine ideas from satisfiability, Markov chain Monte Carlo, belief propagation, and resolution. Learning algorithms make use of conditional likelihood, convex optimization, and inductive logic programming. Markov logic has been successfully applied to problems in information extraction and integration, natural language processing, robot mapping, social networks, computational biology, and others, and is the basis of the open-source Alchemy system. Table of Contents: Introduction / Markov Logic / Inference / Learning / Extensions / Applications / Conclusion

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.

Answer Set Solving in Practice

Answer Set Programming (ASP) is a declarative problem solving approach, initially tailored to modeling problems in the area of Knowledge Representation and Reasoning (KRR). More recently, its attractive combination of a rich yet simple modeling language with high-performance solving capacities has sparked interest in many other areas even beyond KRR. This book presents a practical introduction to ASP, aiming at using ASP languages and systems for solving application problems. Starting from the essential formal foundations, it introduces ASP's solving technology, modeling language and methodology, while illustrating the overall solving process by practical examples. Table of Contents: List of Figures / List of Tables / Motivation / Introduction / Basic modeling / Grounding / Characterizations / Solving / Systems / Advanced modeling / Conclusions

Adversarial Machine Learning

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 malicius 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.

Reasoning with Probabilistic and Deterministic Graphical Models

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. This book provides 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. The new edition includes the notion of influence diagrams, which focus on sequential decision making under uncertainty. We believe the principles outlined in the book 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.

Positive Unlabeled Learning

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.

Learning and Decision-Making from Rank Data

The ubiquitous challenge of learning and decision-making from rank data arises in situations where intelligent systems collect preference and behavior data from humans, learn from the data, and then use the data to help humans make efficient, effective, and timely decisions. Often, such data are represented by rankings. This book surveys some recent progress toward addressing the challenge from the considerations of statistics, computation, and socio-economics. We will cover classical statistical models for rank data, including random utility models, distance-based models, and mixture models. We will discuss and compare classical and state-of-the-art algorithms, such as algorithms based on Minorize-Majorization (MM), Expectation-Maximization (EM), Generalized Method-of-Moments (GMM), rank breaking, and tensor decomposition. We will also introduce principled Bayesian preference elicitation frameworks for collecting rank data. Finally, we will examine socio-economic aspects of statistically desirable decision-making mechanisms, such as Bayesian estimators. This book can be useful in three ways: (1) for theoreticians in statistics and machine learning to better understand the considerations and caveats of learning from rank data, compared to learning from other types of data, especially cardinal data; (2) for practitioners to apply algorithms covered by the book for sampling, learning, and aggregation; and (3) as a textbook for graduate students or advanced undergraduate students to learn about the field. This book requires that the reader has basic knowledge in probability, statistics, and algorithms. Knowledge in social choice would also help but is not required.

Explainable Human-AI Interaction

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.

Statistical Relational Artificial Intelligence

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.

Reasoning with Probabilistic and Deterministic Graphical Models

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.

Action Programming Languages

Artificial systems that think and behave intelligently are one of the most exciting and challenging goals of Artificial Intelligence. Action Programming is the art and science of devising high-level control strategies for autonomous systems which employ a mental model of their environment and which reason about their actions as a means to achieve their goals. Applications of this programming paradigm include autonomous software agents, mobile robots with high-level reasoning capabilities, and General Game Playing. These lecture notes give an in-depth introduction to the current state-of-the-art in action programming. The main topics are knowledge representation for actions, procedural action programming, planning, agent logic programs, and reactive, behavior-based agents. The only prerequisite for understanding the material in these lecture notes is some general programming experience and basic knowledge of classical first-order logic. Table of Contents: Introduction / Mathematical Preliminaries / Procedural Action Programs / Action Programs and Planning / Declarative Action Programs / Reactive Action Programs / Suggested Further Reading

Representation Discovery using Harmonic Analysis

Representations are at the heart of artificial intelligence (AI). This book is devoted to the problem of representation discovery: how can an intelligent system construct representations from its experience? Representation discovery re-parameterizes the state space - prior to the application of information retrieval, machine learning, or optimization techniques - facilitating later inference processes by constructing new task-specific bases adapted to the state space geometry. This book presents a general approach to representation discovery using the framework of harmonic analysis, in particular Fourier and wavelet analysis. Biometric compression methods, the compact disc, the computerized axial tomography (CAT) scanner in medicine, JPEG compression, and spectral analysis of time-series data are among the many applications of classical Fourier and wavelet analysis. A central goal of this book is to show that these analytical tools can be generalized from their usual setting in (infinite-dimensional) Euclidean spaces to discrete (finite-dimensional) spaces typically studied in many subfields of AI. Generalizing harmonic analysis to discrete spaces poses many challenges: a discrete representation of the space must be adaptively acquired; basis functions are not pre-defined, but rather must be constructed. Algorithms for efficiently computing and representing bases require dealing with the curse of dimensionality. However, the benefits can outweigh the costs, since the extracted basis functions outperform parametric bases as they often reflect the irregular shape of a particular state space. Case studies from computer graphics, information retrieval, machine learning, and state space planning are used to illustrate the benefits of the proposed framework, and the challenges that remain to be addressed. Representation discovery is an actively developing field, and the author hopes this book will encourage other researchers to explore this exciting area of research. Table of Contents: Overview / Vector Spaces / Fourier Bases on Graphs / Multiscale Bases on Graphs / Scaling to Large Spaces / Case Study: State-Space Planning / Case Study: Computer Graphics / Case Study: Natural Language / Future Directions

Robot Learning from Human Teachers

Learning from Demonstration (LfD) explores techniques for learning a task policy from examples provided by a human teacher. The field of LfD has grown into an extensive body of literature over the past 30 years, with a wide variety of approaches for encoding human demonstrations and modeling skills and tasks. Additionally, we have recently seen a focus on gathering data from non-expert human teachers (i.e., domain experts but not robotics experts). In this book, we provide an introduction to the field with a focus on the unique technical challenges associated with designing robots that learn from naive human teachers. We begin, in the introduction, with a unification of the various terminology seen in the literature as well as an outline of the design choices one has in designing an LfD system. Chapter 2 gives a brief survey of the psychology literature that provides insights from human social learning that are relevant to designing robotic social learners. Chapter 3 walks through an LfD interaction, surveying the design choices one makes and state of the art approaches in prior work. First, is the choice of input, how the human teacher interacts with the robot to provide demonstrations. Next, is the choice of modeling technique. Currently, there is a dichotomy in the field between approaches that model low-level motor skills and those that model high-level tasks composed of primitive actions. We devote a chapter to each of these. Chapter 7 is devoted to interactive and active learning approaches that allow the robot to refine an existing task model. And finally, Chapter 8 provides best practices for evaluation of LfD systems, with a focus on how to approach experiments with human subjects in this domain.

Network Embedding

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.

Trading Agents

Automated trading in electronic markets is one of the most common and consequential applications of autonomous software agents. Design of effective trading strategies requires thorough understanding of how market mechanisms operate, and appreciation of strategic issues that commonly manifest in trading scenarios. Drawing on research in auction theory and artificial intelligence, this book presents core principles of strategic reasoning that apply to market situations. The author illustrates trading strategy choices through examples of concrete market environments, such as eBay, as well as abstract market models defined by configurations of auctions and traders. Techniques for addressing these choices constitute essential building blocks for the design of trading strategies for rich market applications. The lecture assumes no prior background in game theory or auction theory, or artificial intelligence. Table of Contents: Introduction / Example: Bidding on eBay / Auction Fundamentals / Continuous Double Auctions / Interdependent Markets / Conclusion

Predicting Human Decision-Making

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.

Visual Object Recognition

The visual recognition problem is central to computer vision research. From robotics to information retrieval, many desired applications demand the ability to identify and localize categories, places, and objects. This tutorial overviews computer vision algorithms for visual object recognition and image classification. We introduce primary representations and learning approaches, with an emphasis on recent advances in the field. The target audience consists of researchers or students working in AI, robotics, or vision who would like to understand what methods and representations are available for these problems. This lecture summarizes what is and isn't possible to do reliably today, and overviews key concepts that could be employed in systems requiring visual categorization. Table of Contents: Introduction / Overview: Recognition of Specific Objects / Local Features: Detection and Description / Matching Local Features / Geometric Verification of Matched

Features / Example Systems: Specific-Object Recognition / Overview: Recognition of Generic Object Categories / Representations for Object Categories / Generic Object Detection: Finding and Scoring Candidates / Learning Generic Object Category Models / Example Systems: Generic Object Recognition / Other Considerations and Current Challenges / Conclusions

Trading Agents

Automated trading in electronic markets is one of the most common and consequential applications of autonomous software agents. Design of effective trading strategies requires thorough understanding of how market mechanisms operate, and appreciation of strategic issues that commonly manifest in trading scenarios. Drawing on research in auction theory and artificial intelligence, this book presents core principles of strategic reasoning that apply to market situations. The author illustrates trading strategy choices through examples of concrete market environments, such as eBay, as well as abstract market models defined by configurations of auctions and traders. Techniques for addressing these choices constitute essential building blocks for the design of trading strategies for rich market applications.The lecture assumes no prior background in game theory or auction theory, or artificial intelligence.Table of Contents: Introduction / Example: Bidding on eBay / Auction Fundamentals / Continuous Double Auctions / Interdependent Markets / Conclusion

Multi-Objective Decision Making

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.

valuing people moving forward togetherthe governments annual report on learning disability house of commons
motorola gp338 e user manual