

## Statistical Relational Artificial Intelligence Logic Probability And Computation Synthesis Lectures On Artificial Intelligence And Machine Learning

Advanced statistical modeling and knowledge representation techniques for a newly emerging area of machine learning and probabilistic reasoning; includes introductory material, tutorials for different proposed approaches, and applications. Handling inherent uncertainty and exploiting compositional structure are fundamental to understanding and designing large-scale systems. Statistical relational learning builds on ideas from probability theory and statistics to address uncertainty while incorporating tools from logic, databases and programming languages to represent structure. In Introduction to Statistical Relational Learning, leading researchers in this emerging area of machine learning describe current formalisms, models, and algorithms that enable effective and robust reasoning about richly structured systems and data. The early chapters provide tutorials for material used in later chapters, offering introductions to representation, inference and learning in graphical models, and logic. The book then describes object-oriented approaches, including probabilistic relational models, relational Markov networks, and probabilistic entity-relationship models as well as logic-based formalisms including Bayesian logic programs, Markov logic, and stochastic logic programs. Later chapters discuss such topics as probabilistic models with unknown objects, relational dependency networks, reinforcement learning in relational domains, and information extraction. By presenting a variety of approaches, the book highlights commonalities and clarifies important differences among proposed approaches and, along the way, identifies important representational and algorithmic issues. Numerous applications are provided throughout.

The question, how to combine probability and logic with learning, is getting an increased attention in several disciplines such as knowledge representation, reasoning about uncertainty, data mining, and machine learning simultaneously. This results in the newly emerging subfield known under the names of statistical relational learning and probabilistic inductive logic programming. This book provides an introduction to the field with an emphasis on the methods based on logic programming principles. It is concerned with formalisms and systems, implementations and applications, as well as with the theory of probabilistic inductive logic programming. The 13 chapters of this state-of-the-art survey start with an introduction to probabilistic inductive logic programming; moreover the book presents a detailed overview of the most important probabilistic logic learning formalisms and systems such as relational sequence learning techniques, using kernels with logical representations, Markov logic, the PRISM system, CLP(BN), Bayesian logic programs, and the independent choice logic. The third part provides a detailed account of some show-case applications of probabilistic inductive logic programming. The final part touches upon some theoretical investigations and includes chapters on behavioural comparison of probabilistic logic programming representations and a model-theoretic expressivity analysis.

This English edition monograph is developed and updated from China's best-selling, and award-winning, book on Artificial Intelligence (AI). It covers the foundations as well as the latest developments of AI in a comprehensive and systematic manner. It is a valuable guide for students and researchers on artificial intelligence. A wide range of topics in AI are covered in this book with four distinct features. First of all, the book comprises a comprehensive system, covering the core technology of AI, including the basic theories and techniques of 'traditional' artificial intelligence, and the basic principles and methods of computational intelligence. Secondly, the book focuses on innovation, covering advanced learning methods for machine learning and deep learning techniques and other artificial intelligence that have been widely used in recent years. Thirdly, the theory and practice of the book are highly integrated. There are theories, techniques and methods, as well as many application examples, which will help readers to understand the artificial intelligence theory and its application development. Fourthly, the content structure of the book is quite characteristic, consisting of three parts: (i) knowledge-based artificial intelligence, (ii) data-based artificial intelligence, and (iii) artificial intelligence applications. It is closely related to the core elements of artificial intelligence, namely knowledge, data, algorithms, and computing powers. This reflects the authors' deep understanding of the artificial intelligence discipline.

This comprehensive encyclopedia, in A-Z format, provides easy access to relevant information for those seeking entry into any aspect within the broad field of Machine Learning. Most of the entries in this preeminent work include useful literature references. Artificial Intelligence presents a practical guide to AI, including agents, machine learning and problem-solving simple and complex domains.

This book provides a systematic and comprehensive description of Non-Axiomatic Logic, which is the result of the author's research for about three decades. Non-Axiomatic Logic is designed to provide a uniform logical foundation for Artificial Intelligence, as well as an abstract description of the OC laws of thought OCO followed by the human mind. Different from OC mathematical OCO logic, where the focus is the regularity required when demonstrating mathematical conclusions, Non-Axiomatic Logic is an attempt to return to the original aim of logic, that is, to formulate the regularity in actual human thinking. To achieve this goal, the logic is designed under the assumption that the system has insufficient knowledge and resources with respect to the problems to be solved, so that the OC logical conclusions OCO are only valid with respect to the available knowledge and resources. Reasoning processes according to this logic covers cognitive functions like learning, planning, decision making, problem solving, This book is written for researchers and students in Artificial Intelligence and Cognitive Science, and can be used as a textbook for courses at graduate level, or upper-level undergraduate, on Non-Axiomatic Logic."

This book constitutes the thoroughly refereed post-conference proceedings of the 26th International Conference on Inductive Logic Programming, ILP 2016, held in London, UK, in September 2016. The 10 full papers presented were carefully reviewed and selected from 29 submissions. The papers represent well the current breath of ILP research topics such as predicate invention; graph-based learning; spatial learning; logical foundations; statistical relational learning; probabilistic ILP; implementation and scalability; applications in robotics, cyber security and games.

The multi-volume set LNAI 12975 until 12979 constitutes the refereed proceedings of the European Conference on Machine Learning and Knowledge Discovery in Databases, ECML PKDD 2021, which was held during September 13-17, 2021. The conference was originally planned to take place in Bilbao, Spain, but changed to an online event due to the COVID-19 pandemic. The 210 full papers presented in these proceedings were carefully reviewed and selected from a total of 869 submissions. The volumes are organized in topical sections as follows: Research Track: Part I: Online learning; reinforcement learning; time series, streams, and sequence models; transfer and multi-task learning; semi-supervised and few-shot learning; learning algorithms and applications. Part II: Generative models; algorithms and learning theory; graphs and networks; interpretation, explainability, transparency, safety. Part III: Generative models; search and optimization; supervised learning; text mining and natural language

processing; image processing, computer vision and visual analytics. Applied Data Science Track: Part IV: Anomaly detection and malware; spatio-temporal data; e-commerce and finance; healthcare and medical applications (including Covid); mobility and transportation. Part V: Automating machine learning, optimization, and feature engineering; machine learning based simulations and knowledge discovery; recommender systems and behavior modeling; natural language processing; remote sensing, image and video processing; social media.

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.

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.

Statistical Relational Artificial Intelligence Logic, Probability, and Computation Morgan & Claypool Publishers

Probabilistic information has many uses in an intelligent system. This book explores logical formalisms for representing and reasoning with probabilistic information that will be of particular value to researchers in nonmonotonic reasoning, applications of probabilities, and knowledge representation. It demonstrates that probabilities are not limited to particular applications, like expert systems; they have an important role to play in the formal design and specification of intelligent systems in general. Fahiem Bacchus focuses on two distinct notions of probabilities: one propositional, involving degrees of belief, the other proportional, involving statistics. He constructs distinct logics with different semantics for each type of probability that are a significant advance in the formal tools available for representing and reasoning with probabilities. These logics can represent an extensive variety of qualitative assertions, eliminating requirements for exact point-valued probabilities, and they can represent firstshy;order logical information. The logics also have proof theories which give a formal specification for a class of reasoning that subsumes and integrates most of the probabilistic reasoning schemes so far developed in AI. Using the new logical tools to connect statistical with propositional probability, Bacchus also proposes a system of direct inference in which degrees of belief can be inferred from statistical knowledge and demonstrates how this mechanism can be applied to yield a powerful and intuitively satisfying system of defeasible or default reasoning. Fahiem Bacchus is Assistant Professor of Computer Science at the University of Waterloo, Ontario. Contents: Introduction. Propositional Probabilities. Statistical Probabilities. Combining Statistical and Propositional Probabilities Default Inferences from Statistical Knowledge.

This book constitutes the refereed proceedings of the 15th International Conference of the Italian Association for Artificial Intelligence, AI\*IA 2016, held in Genova, Italy, in November/December 2016. The 39 full papers presented were carefully reviewed and selected from 53 submissions. The papers are organized in topical sections on optimization and evolutionary algorithms; classification, pattern recognition, and computer vision; multi-agent systems; machine learning; semantic web and description logics; natural language processing; planning and scheduling; and formal verification. This book provides an introduction to probabilistic inductive logic programming. It places emphasis on the methods based on logic programming principles and covers formalisms and systems, implementations and applications, as well as theory.

This book is the most comprehensive description of the decades-long Non-Axiomatic Reasoning System (NARS) project, including its philosophical foundation, methodological consideration, conceptual design details, implications in the related fields, and its similarities and differences to many related works in cognitive science. While most current works in Artificial Intelligence (AI) focus on individual aspects of intelligence and cognition, NARS is designed and developed to attack the AI problem as a whole.

Artificial Intelligence is concerned with producing devices that help or replace human beings in their daily activities. Neural-symbolic learning systems play a central role in this task by combining, and trying to benefit from, the advantages of both the neural and symbolic paradigms of artificial intelligence. This book provides a comprehensive introduction to the field of neural-symbolic learning systems, and an invaluable overview of the latest research issues in this area. It is divided into three sections, covering the main topics of neural-symbolic integration - theoretical advances in knowledge representation and learning, knowledge extraction from trained neural networks, and inconsistency handling in neural-symbolic systems. Each section provides a balance of theory and practice, giving the results of applications using real-world problems in areas such as DNA sequence analysis, power systems fault diagnosis, and software requirements

specifications. Neural-Symbolic Learning Systems will be invaluable reading for researchers and graduate students in Engineering, Computing Science, Artificial Intelligence, Machine Learning and Neurocomputing. It will also be of interest to Intelligent Systems practitioners and anyone interested in applications of hybrid artificial intelligence systems. This book constitutes the thoroughly refereed post-proceedings of the 23rd International Conference on Inductive Logic Programming, ILP 2013, held in Rio de Janeiro, Brazil, in August 2013. The 9 revised extended papers were carefully reviewed and selected from 42 submissions. The conference now focuses on all aspects of learning in logic, multi-relational learning and data mining, statistical relational learning, graph and tree mining, relational reinforcement learning, and other forms of learning from structured data.

As the first book devoted to relational data mining, this coherently written multi-author monograph provides a thorough introduction and systematic overview of the area. The first part introduces the reader to the basics and principles of classical knowledge discovery in databases and inductive logic programming; subsequent chapters by leading experts assess the techniques in relational data mining in a principled and comprehensive way; finally, three chapters deal with advanced applications in various fields and refer the reader to resources for relational data mining. This book will become a valuable source of reference for R&D professionals active in relational data mining. Students as well as IT professionals and ambitious practitioners interested in learning about relational data mining will appreciate the book as a useful text and gentle introduction to this exciting new field.

Artificial intelligence (AI) is a field within computer science that is attempting to build enhanced intelligence into computer systems. This book traces the history of the subject, from the early dreams of eighteenth-century (and earlier) pioneers to the more successful work of today's AI engineers. AI is becoming more and more a part of everyone's life. The technology is already embedded in face-recognizing cameras, speech-recognition software, Internet search engines, and health-care robots, among other applications. The book's many diagrams and easy-to-understand descriptions of AI programs will help the casual reader gain an understanding of how these and other AI systems actually work. Its thorough (but unobtrusive) end-of-chapter notes containing citations to important source materials will be of great use to AI scholars and researchers. This book promises to be the definitive history of a field that has captivated the imaginations of scientists, philosophers, and writers for centuries.

Recent advances in the area of lifted inference, which exploits the structure inherent in relational probabilistic models. Statistical relational AI (StaRAI) studies the integration of reasoning under uncertainty with reasoning about individuals and relations. The representations used are often called relational probabilistic models. Lifted inference is about how to exploit the structure inherent in relational probabilistic models, either in the way they are expressed or by extracting structure from observations. This book covers recent significant advances in the area of lifted inference, providing a unifying introduction to this very active field. After providing necessary background on probabilistic graphical models, relational probabilistic models, and learning inside these models, the book turns to lifted inference, first covering exact inference and then approximate inference. In addition, the book considers the theory of liftability and acting in relational domains, which allows the connection of learning and reasoning in relational domains.

Artificial intelligence, or AI, is a cross-disciplinary approach to understanding, modeling, and creating intelligence of various forms. It is a critical branch of cognitive science, and its influence is increasingly being felt in other areas, including the humanities. AI applications are transforming the way we interact with each other and with our environment, and work in artificially modeling intelligence is offering new insights into the human mind and revealing new forms mentality can take. This volume of original essays presents the state of the art in AI, surveying the foundations of the discipline, major theories of mental architecture, the principal areas of research, and extensions of AI such as artificial life. With a focus on theory rather than technical and applied issues, the volume will be valuable not only to people working in AI, but also to those in other disciplines wanting an authoritative and up-to-date introduction to the field. Intended both as a text for advanced undergraduates and graduate students, and as a key reference work for AI researchers and developers, Logical Foundations of Artificial Intelligence is a lucid, rigorous, and comprehensive account of the fundamentals of artificial intelligence from the standpoint of logic. The first section of the book introduces the logicist approach to AI--discussing the representation of declarative knowledge and featuring an introduction to the process of conceptualization, the syntax and semantics of predicate calculus, and the basics of other declarative representations such as frames and semantic nets. This section also provides a simple but powerful inference procedure, resolution, and shows how it can be used in a reasoning system. The next several chapters discuss nonmonotonic reasoning, induction, and reasoning under uncertainty, broadening the logical approach to deal with the inadequacies of strict logical deduction. The third section introduces modal operators that facilitate representing and reasoning about knowledge. This section also develops the process of writing predicate calculus sentences to the metalevel--to permit sentences about sentences and about reasoning processes. The final three chapters discuss the representation of knowledge about states and actions, planning, and intelligent system architecture. End-of-chapter bibliographic and historical comments provide background and point to other works of interest and research. Each chapter also contains numerous student exercises (with solutions provided in an appendix) to reinforce concepts and challenge the learner. A bibliography and index complete this comprehensive work.

The three volume proceedings LNAI 10534 – 10536 constitutes the refereed proceedings of the European Conference on Machine Learning and Knowledge Discovery in Databases, ECML PKDD 2017, held in Skopje, Macedonia, in September 2017. The total of 101 regular papers presented in part I and part II was carefully reviewed and selected from 364 submissions; there are 47 papers in the applied data science, nectar and demo track. The contributions were organized in topical sections named as follows: Part I: anomaly detection; computer vision; ensembles and meta learning; feature selection and extraction; kernel methods; learning and optimization, matrix and tensor factorization; networks and graphs;

neural networks and deep learning. Part II: pattern and sequence mining; privacy and security; probabilistic models and methods; recommendation; regression; reinforcement learning; subgroup discovery; time series and streams; transfer and multi-task learning; unsupervised and semisupervised learning. Part III: applied data science track; nectar track; and demo track.

This first textbook on multi-relational data mining and inductive logic programming provides a complete overview of the field. It is self-contained and easily accessible for graduate students and practitioners of data mining and machine learning.

This book constitutes the thoroughly refereed post-conference proceedings of the 27th International Conference on Inductive Logic Programming, ILP 2017, held in Orléans, France, in September 2017. The 12 full papers presented were carefully reviewed and selected from numerous submissions. Inductive Logic Programming (ILP) is a subfield of machine learning, which originally relied on logic programming as a uniform representation language for expressing examples, background knowledge and hypotheses. Due to its strong representation formalism, based on first-order logic, ILP provides an excellent means for multi-relational learning and data mining, and more generally for learning from structured data.

This book constitutes the refereed proceedings of the 14th International Conference on Scalable Uncertainty Management, SUM 2020, which was held in Bozen-Bolzano, Italy, in September 2020. The 12 full, 7 short papers presented in this volume were carefully reviewed and selected from 30 submissions. Besides that, the book also contains 2 abstracts of invited talks, 2 tutorial papers, and 2 PhD track papers. The conference aims to gather researchers with a common interest in managing and analyzing imperfect information from a wide range of fields, such as artificial intelligence and machine learning, databases, information retrieval and data mining, the semantic web and risk analysis. Due to the Corona pandemic SUM 2020 was held as an virtual event.

The practical benefits of computational logic need not be limited to mathematics and computing. As this book shows, ordinary people in their everyday lives can profit from the recent advances that have been developed for artificial intelligence. The book draws upon related developments in various fields from philosophy to psychology and law. It pays special attention to the integration of logic with decision theory, and the use of logic to improve the clarity and coherence of communication in natural languages such as English. This book is essential reading for teachers and researchers who may be out of touch with the latest developments in computational logic. It will also be useful in any undergraduate course that teaches practical thinking, problem solving or communication skills. Its informal presentation makes the book accessible to readers from any background, but optional, more formal, chapters are also included for those who are more technically oriented.

The 18th International Conference on Inductive Logic Programming was held in Prague, September 10–12, 2008. ILP returned to Prague after 11 years, and it is tempting to look at how the topics of interest have evolved during that time. The ILP community clearly continues to cherish its beloved first-order logic representation framework. This is legitimate, as the work presented at ILP 2008 demonstrated that there is still room for both extending established ILP approaches (such as inverse entailment) and exploring novel logic induction frameworks (such as brave induction). Besides the topics lending ILP research its unique focus, we were glad to see in this year's proceedings a good number of papers contributing to areas such as statistical relational learning, graph mining, or the semantic web. To help open ILP to more mainstream research areas, the conference featured three excellent invited talks from the domains of the semantic web (Frank van Harmelen), bioinformatics (Mark Craven) and cognitive sciences (Josh Tenenbaum). We deliberately looked for speakers who are not directly involved in ILP research. We further invited a tutorial on statistical relational learning (Kristian Kersting) to meet the strong demand to have the topic presented from the ILP perspective. Lastly, Stefano Bertolo from the European Commission was invited to give a talk on the ideal niches for ILP in the current EU-supported research on intelligent content and semantics.

Machine Learning: A Constraint-Based Approach provides readers with a refreshing look at the basic models and algorithms of machine learning, with an emphasis on current topics of interest that includes neural networks and kernel machines. The book presents the information in a truly unified manner that is based on the notion of learning from environmental constraints. While regarding symbolic knowledge bases as a collection of constraints, the book draws a path towards a deep integration with machine learning that relies on the idea of adopting multivalued logic formalisms, like in fuzzy systems. A special attention is reserved to deep learning, which nicely fits the constrained-based approach followed in this book. This book presents a simpler unified notion of regularization, which is strictly connected with the parsimony principle, and includes many solved exercises that are classified according to the Donald Knuth ranking of difficulty, which essentially consists of a mix of warm-up exercises that lead to deeper research problems. A software simulator is also included. Presents fundamental machine learning concepts, such as neural networks and kernel machines in a unified manner Provides in-depth coverage of unsupervised and semi-supervised learning Includes a software simulator for kernel machines and learning from constraints that also includes exercises to facilitate learning Contains 250 solved examples and exercises chosen particularly for their progression of difficulty from simple to complex

Provides an integrated introduction to artificial intelligence. Develops AI representation schemes and describes their uses for diverse applications, from autonomous robots to diagnostic assistants to infobots. DLC: Artificial intelligence.

This SpringerBrief addresses the challenges of analyzing multi-relational and noisy data by proposing several Statistical Relational Learning (SRL) methods. These methods combine the expressiveness of first-order logic and the ability of probability theory to handle uncertainty. It provides an overview of the methods and the key assumptions that allow for adaptation to different models and real world applications. The models are highly attractive due to their compactness and comprehensibility but learning their structure is computationally intensive. To combat this problem, the authors review the use of functional gradients for boosting the structure and the parameters of statistical relational models. The algorithms have been applied successfully in several SRL settings and have been adapted to several real problems from Information extraction in text to medical problems. Including both context and well-tested applications, Boosting Statistical Relational Learning from Benchmarks to Data-Driven Medicine is designed for researchers and professionals in machine learning and data mining. Computer engineers or students interested in statistics, data management, or health informatics will also find this brief a valuable resource.

This book constitutes the proceedings of the 15th European Conference on Logics in Artificial Intelligence, JELIA 2016, held in Larnaca, Cyprus, in November 2015. The 32 full papers and 10 short papers included in this volume were carefully reviewed and selected from 88 submissions. The accepted papers span a number of areas within Logics in AI, including: belief revision, answer set programming, argumentation, probabilistic reasoning, handling inconsistencies, temporal logics and planning, description logics, and decidability and complexity results.

A thought-provoking and wide-ranging exploration of machine learning and the race to build computer intelligences as flexible as our own In the world's top research labs and universities, the race is on to invent the ultimate learning algorithm: one capable of discovering any knowledge from data, and doing anything we want, before we even ask. In The Master Algorithm, Pedro Domingos lifts the veil to give us a peek inside the learning machines that power Google, Amazon, and your smartphone. He assembles a blueprint for the future universal learner--the Master Algorithm--and

discusses what it will mean for business, science, and society. If data-ism is today's philosophy, this book is its bible. This book constitutes the refereed conference proceedings of the 29th International Conference on Inductive Logic Programming, ILP 2019, held in Plovdiv, Bulgaria, in September 2019. The 11 papers presented were carefully reviewed and selected from numerous submissions. Inductive Logic Programming (ILP) is a subfield of machine learning, which originally relied on logic programming as a uniform representation language for expressing examples, background knowledge and hypotheses. Due to its strong representation formalism, based on first-order logic, ILP provides an excellent means for multi-relational learning and data mining, and more generally for learning from structured data. Talks about Logic Programming, Uncertainty Reasoning and Machine Learning. This book includes definitions that circumscribe the area formed by extending Inductive Logic Programming to cases annotated with probability values. It investigates the approach of Learning from proofs and the issue of upgrading Fisher Kernels to Relational Fisher Kernels. The three volume proceedings LNAI 11051 – 11053 constitutes the refereed proceedings of the European Conference on Machine Learning and Knowledge Discovery in Databases, ECML PKDD 2018, held in Dublin, Ireland, in September 2018. The total of 131 regular papers presented in part I and part II was carefully reviewed and selected from 535 submissions; there are 52 papers in the applied data science, nectar and demo track. The contributions were organized in topical sections named as follows: Part I: adversarial learning; anomaly and outlier detection; applications; classification; clustering and unsupervised learning; deep learningensemble methods; and evaluation. Part II: graphs; kernel methods; learning paradigms; matrix and tensor analysis; online and active learning; pattern and sequence mining; probabilistic models and statistical methods; recommender systems; and transfer learning. Part III: ADS data science applications; ADS e-commerce; ADS engineering and design; ADS financial and security; ADS health; ADS sensing and positioning; nectar track; and demo track.

[Copyright: 7c0bd8bf0ddd07dee8bd709cc7f97cd6](#)