Advanced Lectures On Machine Learning Ml Summer Schools 2003 Canberra Australia February 2 14 20

This book discusses one of the major applications of artificial intelligence: the use of machine learning to extract useful information from multimodal data. It discusses the optimization methods that help minimize the error in developing patterns and classifications, which further helps improve prediction and decision-making. The book also presents formulations of real-world machine learning problems, and discusses AI solution methodologies as standalone or hybrid approaches. Lastly, it proposes novel metaheuristic methods to solve complex machine learning problems. Featuring valuable insights, the book helps readers explore new avenues leading toward multidisciplinary research discussions.

Designing molecules and materials with desired properties is an important prerequisite for advancing technology in our modern societies. This requires both the ability to calculate accurate microscopic properties, such as energies, forces and electrostatic multipoles of specific configurations, as well as efficient sampling of potential energy surfaces to obtain corresponding macroscopic properties. Tools that can provide this are accurate first-principles calculations rooted in quantum mechanics, and statistical mechanics, respectively. Unfortunately, they come at a high computational cost that prohibits calculations for large systems and long time-scales, thus presenting a severe bottleneck both for searching the vast chemical compound space and the stupendously many dynamical configurations that a molecule can assume. To overcome this challenge, recently there have been increased efforts to accelerate quantum simulations with machine learning (ML). This emerging interdisciplinary community encompasses chemists, material scientists, physicists, mathematicians and computer scientists, joining forces to contribute to the exciting hot topic of progressing machine learning and AI for molecules and materials. The book that has emerged from a series of workshops provides a snapshot of this rapidly developing field. It contains tutorial material explaining the relevant foundations needed in chemistry, physics as well as machine learning to give an easy starting point for interested readers. In addition, a number of research papers defining the current state-of-the-art are included. The book has five parts (Fundamentals, Incorporating Prior Knowledge, Deep Learning of Atomistic Representations, Atomistic Simulations and Discovery and Design), each prefaced by editorial commentary that puts the respective parts into a broader scientific context.

A comprehensive and self-contained introduction to Gaussian processes, which provide a principled, practical, probabilistic approach to learning in kernel machines. Gaussian processes (GPs) provide a principled, practical, probabilistic approach to learning in kernel machines. GPs have received increased attention in the machine-learning community over the past decade, and this book provides a long-needed systematic and unified treatment of theoretical and practical aspects of GPs in machine learning. The treatment is comprehensive and self-contained, targeted at researchers and students in machine learning and applied statistics. The book deals with the supervised-learning problem for both regression and classification, and includes detailed algorithms. A wide variety of covariance (kernel) functions are presented and their properties discussed. Model selection is discussed both from a Bayesian and a classical perspective. Many connections to other well-known techniques from machine learning and statistics are discussed, including support-vector machines, neural networks, splines, regularization networks, relevance vector machines and others. Theoretical issues including learning curves and the PAC-Bayesian framework are treated, and several approximation methods for learning with large datasets are discussed. The book contains illustrative examples and exercises, and code and datasets are available on the Web. Appendixes provide mathematical background and a discussion of Gaussian Markov processes.

Updated with new code, new projects, and new chapters, Machine Learning with TensorFlow, Second Edition gives readers a solid foundation in machine-learning concepts and the TensorFlow library. Summary Updated with new code, new projects, and new chapters, Machine Learning with TensorFlow, Second Edition gives readers a solid foundation in machine-learning concepts and the TensorFlow library. Written by NASA JPL Deputy CTO and Principal Data Scientist Chris Mattmann, all examples are accompanied by downloadable Jupyter Notebooks for a hands-on experience coding TensorFlow with Python. New and revised content expands coverage of core machine learning algorithms, and advancements in neural networks such as VGG-Face facial identification classifiers and deep speech classifiers. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the technology Supercharge your data analysis with machine learning! ML algorithms automatically improve as they process data, so results get better over time. You don't have to be a mathematician to use ML: Tools like Google's TensorFlow library help with complex calculations so you can focus on getting the answers you need. About the book Machine Learning with TensorFlow, Second Edition is a fully revised guide to building machine learning models using Python and TensorFlow. You'll apply core ML concepts to real-world challenges, such as sentiment analysis, text classification, and image recognition. Hands-on examples illustrate neural network techniques for deep speech processing, facial identification, and auto-encoding with CIFAR-10. What's inside Machine Learning with TensorFlow Choosing the best ML approaches Visualizing algorithms with TensorBoard Sharing results with collaborators Running models in Docker About the reader Requires intermediate Python skills and knowledge of general algebraic concepts like vectors and matrices. Examples use the super-stable 1.15.x branch of TensorFlow and TensorFlow 2.x. About the author Chris Mattmann is the Division Manager of the Artificial Intelligence, Analytics, and Innovation Organization at NASA Jet Propulsion Lab. The first edition of this book was written by Nishant Shukla with Kenneth Fricklas. Table of Contents PART 1 - YOUR MACHINE-LEARNING RIG 1 A machine-learning odyssey 2 TensorFlow essentials PART 2 - CORE LEARNING ALGORITHMS 3 Linear regression and beyond 4 Using regression for call-center volume prediction 5 A gentle introduction to classification 6 Sentiment classification: Large movie-review dataset 7 Automatically clustering data 8 Inferring user activity from Android accelerometer data 9 Hidden Markov models 10 Part-of-speech tagging and word-sense disambiguation PART 3 - THE NEURAL NETWORK PARADIGM 11 A peek into autoencoders 12 Applying autoencoders: The CIFAR-10 image dataset 13 Reinforcement learning 14 Convolutional neural networks 15 Building a real-world CNN: VGG-Face ad VGG-Face Lite 16 Recurrent neural networks 17 LSTMs and automatic speech recognition 18 Sequence-to-sequence models for chatbots 19 Utility landscape

Reinforcement Learning, second edition

16th European Conference on Machine Learning, Porto, Portugal, October 3-7, 2005, Proceedings

Introduction to Semi-supervised Learning

An Introduction for Scientists and Engineers

Machine Learning: ECML 2005

Neural Networks and Deep Learning

Machine Learning has become a key enabling technology for many engineering applications and theoretical problems alike. To further discussions and to dis- minate new results, a Summer School was held on February 11–22, 2002 at the Australian National University. The current book contains a collection of the main talks held during those two weeks in February, presented as tutorial chapters on topics such as Boosting, Data Mining, Kernel Methods, Logic, Reinforcement Learning, and Statistical Learning Theory. The papers provide an in-depth overview of these exciting new areas, contain a large set of references, and thereby provide the interested reader with further information to start or to pursue his own research in these directions. Complementary to the book, a recorded video of the presentations during the Summer School can be obtained at http://mlg. anu. edu. au/summer2002 It is our hope that graduate students, lecturers, and researchers alike will ?nd this book useful in learning and teaching Machine Learning, thereby continuing the mission of the Summer School. Canberra, November 2002 Shahar Mendelson Alexander Smola Research School of Information Sciences and Engineering, The Australian National University Thanks and Acknowledgments We gratefully thank all the individuals and organizations responsible for the success of the workshop.

Introduction to Algorithms for Data Mining and Machine Learning introduces the essential ideas behind all key algorithms and techniques for data mining and machine learning, along with optimization techniques. Its strong formal mathematical approach, well selected examples, and practical software recommendations help readers develop confidence in their data modeling skills so they can process and interpret data for classification, clustering, curve-fitting and predictions. Masterfully balancing theory and practice, it is especially useful for those who need relevant, well explained, but not rigorous (proofs based) background theory and clear guidelines for working with big data. Presents an informal, theorem-free approach with concise, compact coverage of all fundamental topics Includes worked examples that help users increase confidence in their understanding of key algorithms, thus encouraging self-study Provides algorithms and techniques that can be implemented in any programming language, with each chapter including notes about relevant software packages

It is common wisdom that gathering a variety of views and inputs improves the process of decision making, and, indeed, underpins a democratic society. Dubbed "ensemble learning" by researchers in computational intelligence and machine learning, it is known to improve a decision system's robustness and accuracy. Now, fresh developments are allowing researchers to unleash the power of ensemble learning in an increasing range of real-world applications. Ensemble learning algorithms such as "boosting" and "random forest" facilitate solutions to key computational issues such as face recognition and are now being applied in areas as diverse as object tracking and bioinformatics. Responding to a shortage of literature dedicated to the topic, this volume offers comprehensive coverage of state-of-the-art ensemble learning techniques, including the random forest skeleton tracking algorithm in the Xbox Kinect sensor, which bypasses the need for game controllers. At once a solid theoretical study and a practical guide, the volume is a windfall for researchers and practitioners alike.

This book introduces basic machine learning concepts and applications for a broad audience that includes students, faculty, and industry practitioners. We begin by describing how machine learning provides capabilities to computers and embedded systems to learn from data. A typical machine learning algorithm involves training, and generally the performance of a machine learning model improves with more training data. Deep learning is a sub-area of machine learning that involves extensive use of layers of artificial neural networks typically trained on massive amounts of data. Machine and deep learning methods are often used in contemporary data science tasks to address the growing data sets and detect, cluster, and classify data patterns. Although machine learning commercial interest has grown relatively recently, the roots of machine learning go back to decades ago. We note that nearly all organizations, including industry, government, defense, and health, are using machine learning to address a variety of needs and applications. The machine learning paradigms presented can be broadly divided into the following three categories: supervised learning, unsupervised learning, and semi-supervised learning. Supervised learning algorithms focus on learning a mapping function, and they are trained with supervision on labeled data. Supervised learning is further sub-divided into classification and regression algorithms. Unsupervised learning typically does not have access to ground truth, and often the goal is to learn or uncover the hidden pattern in the data. Through semi-supervised learning, one can effectively utilize a large volume of unlabeled data and a limited amount of labeled data to improve machine learning model performances. Deep learning and neural networks are also covered in this book. Deep neural networks have attracted a lot of interest during the last ten years due to the availability of graphics processing units (GPU) computational power, big data, and new software platforms. They have strong capabilities in terms of learning complex mapping functions for different types of data. We organize the book as follows. The book starts by introducing concepts in supervised, unsupervised, and semi-supervised learning. Several algorithms and their inner workings are presented within these three categories. We then continue with a brief introduction to artificial neural network algorithms and their properties. In addition, we cover an array of applications and provide extensive bibliography. The book ends with a summary of the key machine learning concepts.

Methods, Systems, Challenges

Lifelong Machine Learning

Active Learning

Optimization in Machine Learning and Applications

Second Edition

Machine Learning Summer School ... : Revised Lectures

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

Machine learning (ML) is changing virtually every aspect of our lives. Today ML algorithms accomplish tasks that until recently only expert humans could perform. As it relates to finance, this is the most exciting time to adopt a disruptive technology that will transform how everyone invests for generations. Readers will learn how to structure Big data in a way that is amenable to ML algorithms; how to conduct research with ML algorithms on that data; how to use supercomputing methods; how to backtest your discoveries while avoiding false positives. The book addresses real-life problems faced by practitioners on a daily basis, and explains scientifically sound solutions using math, supported by code and examples. Readers become active users who can test the proposed solutions in their particular setting. Written by a recognized expert and portfolio manager, this book will equip investment professionals with the groundbreaking tools needed to succeed in modern finance.

The fundamental mathematical tools needed to understand machine learning include linear algebra, analytic geometry, matrix decompositions, vector calculus, optimization, probability and statistics. These topics are traditionally taught in disparate courses, making it hard for data science or computer science students, or professionals, to efficiently learn the mathematics. This self-contained textbook bridges the gap between mathematical and machine learning texts, introducing the mathematical concepts with a minimum of prerequisites. It uses these concepts to derive four central machine learning methods: linear regression, principal component analysis, Gaussian mixture models and support vector machines. For students and others with a mathematical background, these derivations provide a starting point to machine learning texts. For those learning the mathematics for the first time, the methods help build intuition and practical experience with applying mathematical concepts. Every chapter includes worked examples and exercises to test understanding. Programming tutorials are offered on the book's web site.

This book constitutes the refereed proceedings of the 17th European Conference on Machine Learning, ECML 2006, held, jointly with PKDD 2006. The book presents 46 revised full papers and 36 revised short papers together with abstracts of 5 invited talks, carefully reviewed and selected from 564 papers submitted. The papers present a wealth of new results in the area and address all current issues in machine learning.

Essentials of Game Theory

A Unified Framework

17th European Conference on Machine Learning, Berlin, Germany, September 18-22, 2006, Proceedings

Advanced Lectures

Metric Learning

Introduction to Machine Learning and Bioinformatics

Machine Learning has become a key enabling technology for many engineering applications, investigating scientific questions and theoretical problems alike. To stimulate discussions and to disseminate new results, a summer school series was started in February 2002, the documentation of which is published as LNAI 2600. This book presents revised lectures of two subsequent summer schools held in 2003 in Canberra, Australia, and in Tübingen, Germany. The tutorial lectures included are devoted to statistical learning theory, unsupervised learning, Bayesian inference, and applications in pattern recognition; they provide in-depth overviews of exciting new developments and contain a large number of references. Graduate students, lecturers, researchers and professionals alike will find this book a useful resource in learning and teaching machine learning.

The significantly expanded and updated new edition of a widely used text on reinforcement learning, one of the most active research areas in artificial intelligence. Reinforcement learning, one of the most active research areas in artificial intelligence, is a computational approach to learning whereby an agent tries to maximize the total amount of reward it receives while interacting with a complex, uncertain environment. In Reinforcement Learning, Richard Sutton and Andrew Barto provide a clear and simple account of the field's key ideas and algorithms. This second edition has been significantly expanded and updated, presenting new topics and updating coverage of other topics. Like the first edition, this second edition focuses on core online learning algorithms, with the more mathematical material set off in shaded boxes. Part I covers as much of reinforcement learning as possible without going beyond the tabular case for which exact solutions can be found. Many algorithms presented in this part are new to the second edition, including UCB, Expected Sarsa, and Double Learning. Part II extends these ideas to function approximation, with new sections on such topics as artificial neural networks and the Fourier basis, and offers expanded treatment of off-policy learning and policy-gradient methods. Part III has new chapters on reinforcement learning's relationships to psychology and neuroscience, as well as an updated case-studies chapter including AlphaGo and AlphaGo Zero, Atari game playing, and IBM Watson's wagering strategy. The final chapter discusses the future societal impacts of reinforcement learning.

An introduction to a broad range of topics in deep learning, covering mathematical and conceptual background, deep learning techniques used in industry, and research perspectives. "Written by three experts in the field, Deep Learning is the only comprehensive book on the subject." —Elon Musk, cochair of OpenAI; cofounder and CEO of Tesla and SpaceX Deep learning is a form of machine learning that enables computers to learn from experience and understand the world in terms of a hierarchy of concepts. Because the computer gathers knowledge from experience, there is no need for a human computer operator to formally specify all the knowledge that the computer needs. The hierarchy of concepts allows the computer to learn complicated concepts by building them out of simpler ones; a graph of these hierarchies would be many layers deep. This book introduces a broad range of topics in deep learning. The text offers mathematical and conceptual background, covering relevant concepts in linear algebra, probability theory and information theory, numerical computation, and machine learning. It describes deep learning techniques used by practitioners in industry, including deep feedforward networks, regularization, optimization algorithms, convolutional networks, sequence modeling, and practical methodology; and it surveys such applications as natural language processing, speech recognition, computer vision, online recommendation systems, bioinformatics, and videogames. Finally, the book offers research perspectives, covering such theoretical topics as linear factor models, autoencoders, representation learning, structured probabilistic models, Monte Carlo methods, the partition function, approximate inference, and deep generative models. Deep Learning can be used by undergraduate or graduate students planning careers in either industry or research, and by software engineers who want to begin using deep learning in their products or platforms. A website offers supplementary material for both readers and instructors.

This book is the outcome of a series of discussions at the Philips Symposium on Intelligent Algorithms, which was held in Eindhoven on December 2002. It contains many exciting and practical examples from this newly developing research field, which can be positioned at the intersection of computer science, discrete mathematics, and artificial intelligence. The examples include machine learning, content management, vision, speech, content augmentation, profiling, music retrieval, feature extraction, audio and video fingerprinting, resource management, multimedia servers, network scheduling, and IC design.

An AI Perspective

A Concise Multidisciplinary Introduction

Statistical Machine Learning

Machine Learning: ECML 2006

A Textbook

An overview of recent efforts in the machine learning community to deal with dataset and covariate shift, which occurs when test and training inputs and outputs have different distributions. Dataset shift is a common problem in predictive modeling that occurs when the joint distribution of inputs and outputs differs between training and test stages. Covariate shift, a particular case of dataset shift, occurs when only the input distribution changes. Dataset shift is present in most practical applications, for reasons ranging from the bias introduced by experimental design to the irreproducibility of the testing conditions at training time. (An example is -email spam filtering, which may fail to recognize spam that differs in form from the spam the automatic filter has been built on.) Despite this, and despite the attention given to the apparently similar problems of semi-supervised learning and active learning, dataset shift has received relatively little attention in the machine learning community until recently. This volume offers an overview of current efforts to deal with dataset and covariate shift. The chapters offer a mathematical and philosophical introduction to the problem, place dataset shift in relationship to transfer learning, transduction, local learning, active learning, and semi-supervised learning, provide theoretical views of dataset and covariate shift (including decision theoretic and Bayesian perspectives), and present algorithms for covariate shift. Contributors: Shai Ben-David, Steffen Bickel, Karsten Borgwardt, Michael Brückner, David Corfield, Amir Globerson, Arthur Gretton, Lars Kai Hansen, Matthias Hein, Jiayuan Huang, Choon Hui Teo, Takafumi Kanamori, Klaus-Robert Müller, Sam Roweis, Neil Rubens, Tobias Scheffer, Marcel Schmittfull, Bernhard Schölkopf Hidetoshi Shimodaira, Alex Smola, Amos Storkey, Masashi Sugiyama

Accompanying CD-RM contains Machine learning office software, MLO guide (pdf) and examples of data.

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

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.

Machine Learning Summer School 2002, Canberra, Australia, February 11-22, 2002, Revised Lectures

ML Summer Schools 2003, Canberra, Australia, February 2-14, 2003, Tübingen, Germany, August 4-16, 2003, Revised Lectures

Methods and Applications

From Theory to Algorithms

Ensemble Machine Learning

Machine Learning Meets Quantum Physics

An introduction to key concepts and techniques in probabilistic machine learning for civil engineering students and professionals; with many step-by-step examples, illustrations, and exercises. This book introduces probabilistic machine learning concepts to civil engineering students and professionals, presenting key approaches and techniques in a way that is accessible to readers without a specialized background in statistics or computer science. It presents different methods clearly and directly, through step-by-step examples, illustrations, and exercises. Having mastered the material, readers will be able to understand the more advanced machine learning literature from which this book draws. The book presents key approaches in the three subfields of probabilistic machine learning: supervised learning, unsupervised learning, and reinforcement learning. It first covers the background knowledge required to understand machine learning, including linear algebra and probability theory. It goes on to present Bayesian estimation, which is behind the formulation of both supervised and unsupervised learning methods, and Markov chain Monte Carlo methods, which enable Bayesian estimation in certain complex cases. The book then covers approaches associated with supervised learning, including regression methods and classification methods, and notions associated with unsupervised learning, including clustering, dimensionality reduction, Bayesian networks, state-space models, and model calibration. Finally, the book introduces fundamental concepts of rational decisions in uncertain contexts and rational decision-making in uncertain and sequential contexts. Building on this, the book describes the basics of reinforcement learning, whereby a virtual agent learns how to make optimal decisions through trial and error while interacting with its environment.

This modern and self-contained book offers a clear and accessible introduction to the important topic of machine learning with neural networks. In addition to describing the mathematical principles of the topic, and its historical evolution, strong connections are drawn with underlying methods from statistical physics and current applications within science and engineering. Closely based around a well-established undergraduate course, this pedagogical text provides a solid understanding of the key aspects of modern machine learning with artificial neural networks, for students in physics, mathematics, and engineering. Numerous exercises expand and reinforce key concepts within the book and allow students to hone their programming skills. Frequent references to current research develop a detailed perspective on the state-of-the-art in machine learning research.

This book covers both classical and modern models in deep learning. The primary focus is on the theory and algorithms of deep learning. The theory and algorithms of neural networks are particularly important for understanding important concepts, so that one can understand the important design concepts of neural architectures in different applications. Why do neural networks work? When do they work better than off-the-shelf machine-learning models? When is depth useful? Why is training neural networks so hard? What are the pitfalls? The book is also rich in discussing different applications in order to give the practitioner a flavor of how neural architectures are designed for different types of problems. Applications associated with many different areas like recommender systems, machine translation, image captioning, image classification, reinforcement-learning based gaming, and text analytics are covered. The chapters of this book span three categories: The basics of neural networks: Many traditional machine learning models can be understood as special cases of neural networks. An emphasis is placed in the first two chapters on understanding the relationship between traditional machine learning and neural networks. Support vector machines, linear/logistic regression, singular value decomposition, matrix factorization, and recommender systems are shown to be special cases of neural networks. These methods are studied together with recent feature engineering methods like word2vec. Fundamentals of neural networks: A detailed discussion of training and regularization is provided in Chapters 3 and 4. Chapters 5 and 6 present radial-basis function (RBF) networks and restricted Boltzmann machines. Advanced topics in neural networks: Chapters 7 and 8 discuss recurrent neural networks and convolutional neural networks. Several advanced topics like deep reinforcement learning, neural Turing machines, Kohonen self-organizing maps, and generative adversarial networks are introduced in Chapters 9 and 10. The book is written for graduate students, researchers, and practitioners. Numerous exercises are available along with a solution manual to aid in classroom teaching. Where possible, an application-centric view is highlighted in order to provide an understanding of the practical uses of each class of techniques.

Introduces machine learning and its algorithmic paradigms, explaining the principles behind automated learning approaches and the considerations underlying their usage.

Mathematics for Machine Learning

An Introduction

Deep Learning

Advanced Lectures on Machine Learning

Algorithms in Ambient Intelligence

Understanding Machine Learning

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.

Machine learning allows computers to learn and discern patterns without actually being programmed. When Statistical techniques and machine learning are combined together they are a powerful tool for analysing various kinds of data in many computer science/engineering areas including, image processing, speech processing, natural language processing, robot control, as well as in fundamental sciences such as biology, medicine, astronomy, physics, and materials. Introduction to Statistical Machine Learning provides a general introduction to machine learning that covers a wide range of topics concisely and will help you bridge the gap between theory and practice. Part I discusses the fundamental concepts of statistics and probability that are used in describing machine learning algorithms. Part II and Part III explain the two major approaches of machine learning techniques; generative methods and discriminative methods. While Part III provides an in-depth look at advanced topics that play essential roles in making machine learning algorithms more useful in practice. The accompanying MATLAB/Octave programs provide you with the necessary practical skills needed to accomplish a wide range of data analysis tasks. Provides the necessary background material to understand machine learning such as statistics, probability, linear algebra, and calculus. Complete coverage of the generative approach to statistical pattern recognition and the discriminative approach to statistical machine learning. Includes MATLAB/Octave programs so that readers can test the algorithms numerically and acquire both mathematical and practical skills in a wide range of data analysis tasks Discusses a wide range of applications in machine learning and statistics and provides examples drawn from image processing, speech processing, natural language processing, robot control, as well as biology, medicine, astronomy, physics, and materials.

This open access book presents the first comprehensive overview of general methods in Automated Machine Learning (AutoML), collects descriptions of existing systems based on these methods, and discusses the first series of international challenges of AutoML systems. The recent success of commercial ML applications and the rapid growth of the field has created a high demand for off-the-shelf ML methods that can be used easily and without expert knowledge. However, many of the recent machine learning successes crucially rely on human experts, who manually select appropriate ML architectures (deep learning architectures or more traditional ML workflows) and their hyperparameters. To overcome this problem, the field of AutoML targets a progressive automation of machine learning, based on principles from optimization and machine learning itself. This book serves as a point of entry into this quickly-developing field for researchers and advanced students alike, as well as providing a reference for practitioners aiming to use AutoML in their work.

Advanced Lectures on Machine LearningMachine Learning Summer School 2002, Canberra, Australia, February 11-22, 2002, Revised LecturesSpringer

Introduction to Statistical Machine Learning

Machine Learning with TensorFlow, Second Edition

Advances in Financial Machine Learning

Probabilistic Machine Learning for Civil Engineers

Machine Learning and Its Applications

Machine and Deep Learning Algorithms and Applications

A new edition of a graduate-level machine learning textbook that focuses on the analysis and theory of algorithms. This book is a general introduction to machine learning that can serve as a textbook for graduate students and a reference for researchers. It covers fundamental modern topics in machine learning while providing the theoretical basis and conceptual tools needed for the discussion and justification of algorithms. It also describes several key aspects of the application of these algorithms. The authors aim to present novel theoretical tools and concepts while giving concise proofs even for relatively advanced topics. Foundations of Machine Learning is unique in its focus on the analysis and theory of algorithms. The first four chapters lay the theoretical foundation for what follows; subsequent chapters are mostly self-contained. Topics covered include the Probably Approximately Correct (PAC) learning framework; generalization bounds based on Rademacher complexity and VC-dimension; Support Vector Machines (SVMs); kernel methods; boosting; on-line learning; multi-class classification; ranking; regression; algorithmic stability; dimensionality reduction; learning automata and languages; and reinforcement learning. Each chapter ends with a set of exercises. Appendixes provide additional material including concise probability review. This second edition offers three new chapters, on model selection, maximum entropy models, and conditional entropy models. New material in the appendixes includes a major section on Fenchel duality, expanded coverage of concentration inequalities, and an entirely new entry on information theory. More than half of the exercises are new to this edition.

An up-to-date account of the interplay between optimization and machine learning, accessible to students and researchers in both communities. The interplay between optimization and machine learning is one of the most important developments in modern computational science. Optimization formulations and methods are proving to be vital in designing algorithms to extract essential knowledge from huge volumes of data. Machine learning, however, is not simply a consumer of optimization technology but a rapidly evolving field that is itself generating new optimization ideas. This book captures the state of the art of the interaction between optimization and machine learning in a way that is accessible to researchers in both fields. Optimization approaches have enjoyed prominence in machine learning because of their wide applicability and attractive theoretical properties. The increasing complexity, size, and variety of today's machine learning models call for the reassessment of existing assumptions. This book starts the process of reassessment. It describes the resurgence in novel contexts of established frameworks such as first-order methods, stochastic approximations, convex relaxations, interior-point methods, and proximal methods. It also devotes attention to newer themes such as regularized optimization, robust optimization, gradient and subgradient methods, splitting techniques, and second-order methods. Many of these techniques draw inspiration from other fields, including operations research, theoretical computer science, and subfields of optimization. The book will enrich the ongoing cross-fertilization between the machine learning community and these other fields, and within the broader optimization community.

This is a technical overview of the field of adversarial machine learning which has emerged to study vulnerabilities of machine learning approaches in adversarial settings and to develop techniques to make learning robust to adversarial manipulation. After reviewing machine learning concepts and approaches, as well as common use cases of these in adversarial settings, we present a general categorization of attacks on machine learning. We then address two major categories of attacks and associated defenses: decision-time attacks, in which an adversary changes the nature of instances seen by a learned model at the time of prediction in order to cause errors, and poisoning or training time attacks, in which the actual training dataset is maliciously modified. In our final chapter devoted to technical content, we discuss recent techniques for attacks on deep learning, as well as approaches for improving robustness of deep neural networks. We conclude with a discussion of several important issues in the area of adversarial learning that in our view warrant further research. The increasing abundance of large high-quality datasets, combined with significant technical advances over the last several decades have made machine learning into a major tool employed across a broad array of tasks including vision, language, finance, and security. However, success has been accompanied with important new challenges: many applications of machine learning are adversarial in nature. Some are adversarial because they are safety critical, such as autonomous driving. An adversary in these applications can be a malicious party aimed at causing congestion or accidents, or may even model unusual situations that expose vulnerabilities in the prediction engine. Other applications are adversarial because their task and/or the data they use are. For example, an important class of problems in security involves detection, such as malware, spam, and intrusion detection. The use of machine learning for detecting malicious entities creates an incentive among adversaries to evade detection by changing their behavior or the content of malicious objects they develop. Given the increasing interest in the area of adversarial machine learning, we hope this book provides readers with the tools necessary to successfully engage in research and practice of machine learning in adversarial settings.

Semi-supervised learning is a learning paradigm concerned with the study of how computers and natural systems such as humans learn in the presence of both labeled and unlabeled data. Traditionally, learning has been studied either in the unsupervised paradigm (e.g., clustering, outlier detection) where all the data are unlabeled, or in the supervised paradigm (e.g., classification, regression) where all the data are labeled. The goal of semi-supervised learning is to understand how combining labeled and unlabeled data may change the learning behavior, and design algorithms that take advantage of such a combination. Semi-supervised learning is of great interest in machine learning and data mining because it can use readily available unlabeled data to improve supervised learning tasks when the labeled data are scarce or expensive. Semi-supervised learning also shows potential as a quantitative tool to understand human category learning, where most of the input is self-evidently unlabeled. In this introductory book, we present some popular semi-supervised learning models, including self-training, mixture models, co-training and multiview learning, graph-based methods, and semi-supervised support vector machines. For each model, we discuss its basic mathematical formulation. The success of semi-supervised learning depends critically on some underlying assumptions. We emphasize the assumptions made by each model and give counterexamples when appropriate to demonstrate the limitations of the different models. In addition, we discuss semi-supervised learning for cognitive psychology. Finally, we give a computational learning theoretic perspective on semi-supervised learning, and we conclude the book with a brief discussion of open questions in the field. Table of Contents: Introduction to Statistical Machine Learning / Overview of Semi-Supervised Learning / Mixture Models and EM / Co-Training / Graph-Based Semi-Supervised Learning / Semi-Supervised Support Vector Machines / Human Semi-Supervised Learning / Theory and Outlook

Planning with Markov Decision Processes

Introduction to Algorithms for Data Mining and Machine Learning

Automated Machine Learning

Theory, Applications, and Software

Optimization for Machine Learning

Machine Learning for Spatial Environmental Data

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

In recent years machine learning has made its way from artificial intelligence into areas of administration, commerce, and industry. Data mining is perhaps the most widely known demonstration of this migration, complemented by less publicized applications of machine learning like adaptive systems in industry, financial prediction, medical diagnosis and the construction of user profiles for Web browsers. This book presents the capabilities of machine learning methods and ideas on how these methods could be used to solve real-world problems. The first ten chapters assess the current state of the art of machine learning, from symbolic concept learning and conceptual clustering to case-based reasoning, neural networks, and genetic algorithms. The second part introduces the reader to innovative applications of ML techniques in fields such as data mining, knowledge discovery, human language technology, user modeling, data analysis, discovery science, agent technology, finance, etc.

The recent rapid growth in the variety and complexity of new machine learning architectures requires the development of improved methods for designing, analyzing, evaluating, and communicating machine learning technologies. Statistical Machine Learning: A Unified Framework provides students, engineers, and scientists with tools from mathematical statistics and nonlinear optimization theory to become experts in the field of machine learning. In particular, the material in this text directly supports the mathematical analysis and design of old, new, and not-yet-invented nonlinear high-dimensional machine learning algorithms. Features: Unified empirical risk minimization framework supports rigorous mathematical analyses of widely used supervised, unsupervised, and reinforcement machine learning algorithms Matrix calculus methods for supporting machine learning analysis and design applications Explicit conditions for ensuring convergence of adaptive, batch, minibatch, MCEM, and MCMC learning algorithms that minimize both unimodal and multimodal objective functions Explicit conditions for characterizing asymptotic properties of M-estimators and model selection criteria such as AIC and BIC in the presence of possible model misspecification This advanced text is suitable for graduate students or highly motivated undergraduate students in statistics, computer science, electrical engineering, and applied mathematics. The text is self-contained and only assumes knowledge of lower-division linear algebra and upper-division probability theory. Students, professional engineers, and multidisciplinary scientists possessing these minimal prerequisites will find this text challenging yet accessible. About the Author: Richard M. Golden (Ph.D., M.S.E.E., B.S.E.E.) is Professor of Cognitive Science and Participating Faculty Member in Electrical Engineering at the University of Texas at Dallas. Dr. Golden has published articles and given talks at scientific conferences on a wide range of topics in the fields of both statistics and machine learning over the past three decades. His long-term research interests include identifying conditions for the convergence of deterministic and stochastic machine learning algorithms and investigating estimation and inference in the presence of possibly misspecified probability models.

Lucidly Integrates Current Activities Focusing on both fundamentals and recent advances, Introduction to Machine Learning and Bioinformatics presents an informative and accessible account of the ways in which these two increasingly intertwined areas relate to each other. Examines Connections between Machine Learning & Bioinformatics The book begins with a brief historical overview of the technological developments in biology. It then describes the main problems in bioinformatics and the fundamental concepts and algorithms of machine learning. After forming this foundation, the authors explore how machine learning techniques apply to bioinformatics problems, such as electron density map interpretation, biclustering, DNA sequence analysis, and tumor classification. They also include exercises at the end of some chapters and offer supplementary materials on their website. Explores How Machine Learning Techniques Can Help Solve Bioinformatics Problems Shedding light on aspects of both machine learning and bioinformatics, this text shows how the innovative tools and techniques of machine learning help extract knowledge from the deluge of information produced by today's biological experiments.

Machine Learning with Neural Networks

Gaussian Processes for Machine Learning

Foundations of Machine Learning, second edition

Adversarial Machine Learning

Dataset Shift in Machine Learning

This book constitutes the refereed proceedings of the 16th European Conference on Machine Learning, ECML 2005, jointly held with PKDD 2005 in Porto, Portugal, in October 2005. The 40 revised full papers and 32 revised short papers presented together with abstracts of 6 invited talks were carefully reviewed and selected from 335 papers submitted to ECML and 30 papers submitted to both, ECML and PKDD. The papers present a wealth of new results in the area and address all current issues in machine learning.