Machine Learning In Non Stationary Environments Introduction To Covariate Shift Adaptation Adaptive Computation And Machine Learning Series

"This book is a vital reference source that trends in data analytics and potential technologies that will facilitate insight in various domains of science, industry, business, and consumer applications. It also explores the latest concepts, algorithms, and techniques of deep learning and data mining and analysis"--

This book constitutes the refereed proceedings of the 5th International Symposium on Cyber Security Cryptography and Machine Learning, CSCML 2021, held in Be'er Sheva, Israel, in July 2021. The 22 full and 13 short papers presented together with a keynote paper in this volume were carefully reviewed and selected from 48 submissions. They deal with the theory, design, analysis, implementation, or application of cyber security, cryptography and machine learning systems and networks, and conceptually innovative topics in these research areas.

This two volume set (LNCS 6791 and LNCS 6792) constitutes the refereed proceedings of the 21th International Conference on Artificial Neural Networks, ICANN 2011, held in Espoo, Finland, in June 2011. The 106 revised full or poster papers presented were carefully reviewed and selected from numerous submissions. ICANN 2011 had two basic tracks: brain-inspired computing and machine learning research, with strong cross-disciplinary interactions and applications.

The proceedings set LNCS 12891, LNCS 12892, LNCS 12893, LNCS 12894 and LNCS 12895 constitute the proceedings of the 30th International Conference on Artificial Neural Networks, ICANN 2021, held in Bratislava, Slovakia, in September 2021.* The total of 265 full papers presented in these proceedings was carefully reviewed and selected from 496 submissions, and organized in 5 volumes. In this volume, the papers focus on topics such as computer vision and object detection, convolutional neural networks and kernel methods, deep learning and optimization, distri continual learning, explainable methods, few-shot learning and generative adversarial networks. *The conference was held online 2021 due to the COVID-19 pandemic.

Intelligent Data Engineering and Automated Learning – IDEAL 2019

Learning in Non-Stationary Environments

Artificial Neural Networks and Machine Learning – ICANN 2021

Artificial Neural Networks and Machine Learning -- ICANN 2012

Machine Learning in Finance

Machine Learning Proceedings 1994

An Introduction

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

 COLOR VERSION New edition updated 2021! Machine learning is one of the most powerful artificial intelligence techniques, capable of efficiently managing and analyzing large amounts of data, to provide accurate predictions, automated decisions and deliver unprecedented business benefits. This volume aims to illustrate in the simplest possible way which are the main approaches in the Machine Learning universe, as well as providing some examples of real applications from which the reader can draw inspiration to understand the benefits and design applications of common interest. Among the covered topics you will find: * Practical applications: regression and classification predictions * Sentiment analysis * Speech Analytics * Image recognition * Performance analysis * Numerous examples and graphical displays of the results * Wavelet Transform for AI non-stationary signal processing * Supervised, Unsupervised and Reinforcement Learning * Artificial Neural Networks, Convolutional Neural Networks, Recurrent Neural Networks * AutoML * MLOps and Pipeline for Distributed Architectures to improve Governance and Scalability The author describes the essential principles and methods of ML (Machine Learning) clearly, making the book suitable even for non-IT readers or data scientists who are experts in the field. Business innovation managers and departments can also benefit from reading this book to better understand how ML can streamline its operations and increase productivity, with an eye to the future. After a first introduction to the concepts of data science and the nomenclature often adopted when it comes to Machine Learning, the book offers a description of the three main methodologies adopted today, trying to analyze both the benefits and the critical issues. Some of the most common learning models are illustrated and the various steps for preparing the data are then analyzed together with the training, testing and accuracy assessment phases. Some of the IT tools that can be used to work on Machine Learning are then described (with emphasis on Open Source ones). The second part of the book deals with different techniques of Regression, Classification and Deep Learning, as well as the methodologies to optimize the results and combine the adopted algorithms. We examine the subject of model interpretability and also of AI security, to move on to an overview of visualization and analysis techniques during Machine Learning processes. The final part focuses on real applications. Two practical cases related to real business applications are dealt with, the approaches to face them, the tools adopted are described and all the source code is made available, commenting it step by step for greater understanding. This volume tries to deal with concepts related to the world of Machine Learning using a language suitable for a wider audience possible, because Machine Learning is part of the fascinating vast world of data science, which brings together various skills: technology, analysis and business understanding.

Theory, algorithms, and applications of machine learning techniques to overcome "covariate shift" non-stationarity. As the power of computing has grown over the past few decades, the field of machine learning has advanced rapidly in both theory and practice. Machine learning methods are usually based on the assumption that the data generation mechanism does not change over time. Yet real-world applications of machine learning, including image recognition, natural language processing, speech recognition, robot control, and bioinformatics, often violate this common assumption. Dealing with non-stationarity is one of modern machine learning's greatest challenges. This book focuses on a specific non-stationary environment known as covariate shift, in which the distributions of inputs (queries) change but the conditional distribution of outputs (answers) is unchanged, and presents machine learning theory, algorithms, and applications to overcome this variety of non-stationarity. After reviewing the state-of-the-art research in the field, the authors discuss topics that include learning under covariate shift, model selection, importance estimation, and active learning. They describe such real world applications of covariate shift adaption as brain-computer interface, speaker identification, and age prediction from facial images. With this book, they aim to encourage future research in machine learning, statistics, and engineering that strives to create truly autonomous learning machines able to learn under non-stationarity.

This Special Issue focused on novel vision-based approaches, mainly related to computer vision and machine learning, for the automatic analysis of human behaviour. We solicited submissions on the following topics: information theory-based pattern classification, biometric recognition, multimodal human analysis, low resolution human activity analysis, face analysis, abnormal behaviour analysis, unsupervised human analysis scenarios, 3D/4D human pose and shape estimation, human analysis in virtual/augmented reality, affective computing, social signal processing, personality computing, activity recognition, human tracking in the wild, and application of information-theoretic concepts for human behaviour analysis. In the end, 15 papers were accepted for this special issue. These papers, that are reviewed in this editorial, analyse human behaviour from the aforementioned perspectives, defining in most of the cases the state of the art in their corresponding field.

Forecasting: principles and practice

Dataset Shift in Machine Learning

Methods and Applications

Complex-valued Neural Networks

European Conference, Antwerp, Belgium, September 15-19, 2008, Proceedings, Part I

Concepts, Methodologies, Tools, and Applications

Introduction to Covariate Shift Adaptation

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.

This book introduces machine learning methods in finance. It presents a unified treatment of machine learning and various statistical and computational disciplines in quantitative finance, such as financial econometrics and discrete time stochastic control, with an emphasis on how theory and hypothesis tests inform the choice of algorithm for financial data modeling and decision making. With the trend towards increasing computational resources and larger datasets, machine learning has grown into an important skillset for the finance industry. This book is written for advanced graduate students and academics in financial econometrics, mathematical finance and applied statistics, in addition to quants and data scientists in the field of quantitative finance. Machine Learning in Finance: From Theory to Practice is divided into three parts, each part covering theory and applications. The first presents supervised learning for cross-sectional data from both a Bayesian and frequentist perspective. The more advanced material places a firm emphasis on neural networks, including deep learning, as well as Gaussian processes, with examples in investment management and derivative modeling. The second part presents supervised learning for time series data, arguably the most common data type used in finance with examples in trading, stochastic volatility and fixed income modeling. Finally, the third part presents reinforcement learning and its applications in trading, investment and wealth management. Python code examples are provided to support the readers' understanding of the methodologies and applications. The book also includes more than 80 mathematical and programming exercises, with worked solutions available to instructors. As a bridge to research in this emergent field, the final chapter presents the frontiers of machine learning in finance from a researcher's perspective, highlighting how many well-known concepts in statistical physics are likely to emerge as important methodologies for machine learning in finance.

Preface v 1 On the History of the Growth-Optimal Portfolio M. M. Christensen 1 2 Empirical Log-Optimal Portfolio Selections: A Survey L. Györfi Gy. Ottucsàak A. Urbán 81 3 Log-Optimal Portfolio-Selection Strategies with Proportional Transaction Costs L. Györfi H. Walk 119 4 Growth-Optimal Portfoho Selection with Short Selling and Leverage M. Horváth A. Urbán 153 5 Nonparametric Sequential Prediction of Stationary Time Series L. Györfi Gy. Ottucsák 179 6 Empirical Pricing American Put Options L. Györfi A. Telcs 227 Index 249

In today's world, we are increasingly exposed to the words 'machine learning' (ML), a term which sounds like a panacea designed to cure all problems ranging from image recognition to machine language translation. Over the past few years, ML has gradually permeated the financial sector, reshaping the landscape of quantitative finance as we know it.An Introduction to Machine Learning in Quantitative Finance aims to demystify ML by uncovering its underlying mathematics and showing how to apply ML methods to real-world financial data. In this book the authorsFeatured with the balance of mathematical theorems and practical code examples of ML, this book will help you acquire an in-depth understanding of ML algorithms as well as hands-on experience. After reading An Introduction to Machine Learning in Quantitative Finance, ML tools will not be a black box to you anymore, and you will feel confident in successfully applying what you have learnt to empirical financial data!

Deployable Machine Learning for Security Defense

22nd International Conference on Artificial Neural Networks, Lausanne, Switzerland, September 11-14, 2012, Proceedings, Part II

Computer Decision-Making

Artificial Neural Networks and Machine Learning - ICANN 2011

Theories and Applications

Statistical Machine Learning

The Sixth Chakra of Artificial Intelligence (color Edition)

Reinforcement learning is a mathematical framework for developing computer agents that can learn an optimal behavior by relating generic reward signals with its past actions. With numerous successful applications in business intelligence, plant control, and gaming, the RL framework is ideal for decision making in unknown environments with large amo

Machine learning is increasingly applied to time series data, as it constitutes an attractive alternative to forecasts based on traditional time series models. For independent and identically distributed observations, cross-validation is the prevalent scheme for estimating out-of-sample performance in both model selection and assessment. For time series data, however, it is unclear whether forwardvalidation schemes, i.e., schemes that keep the temporal order of observations, should be preferred. In this paper, we perform a comprehensive empirical study of eight common validation schemes. We introduce a study design that perturbs global stationarity by introducing a slow evolution of the underlying data-generating process. Our results demonstrate that, even for relatively small perturbations, commonly used cross-validation schemes often yield estimates with the largest bias and variance, and forward-validation schemes yield better estimates of the out-of-sample error. We provide an interpretation of these results in terms of an additional evolution-induced bias and the sample-size dependent estimation error. Using a large-scale financial data set, we demonstrate the practical significance in a replication study of a statistical arbitrage problem. We conclude with some general guidelines on the selection of suitable validation schemes for time series data.

This book attempts to provide a unified overview of the broad field of Machine Learning and its Practical implementation. This book is a survey of the state of art. It breaks this massive subject into comprehensible parts piece by piece. The objective is to focus on basic principles of machine learning with some leading edge topics. This book addresses a full spectrum of machine learning programming. The emphasis is to solve lot many programming examples using step-by step practical implementation of machine learning algorithms. To facilitate easy understanding of machine learning, this book has been written in such a simple style that a student thinks as if a teacher is sitting behind him and guiding him. This book is written as per the new syllabus of different Universities of India. It also Cover the syllabus of B.Tech.(CSE/IT), MCA, BCA of Delhi University, Delhi. GGSIPU, MDU, RGTU, Nagpur University, UTU, APJ Abdul Kalam University so on. The book is intended for both academic and professional audience.

In recent years, complex-valued neural networks have widened the scope of application in optoelectronics, imaging, remote sensing, quantum neural devices and systems, spatiotemporal analysis of physiological neural systems, and artificial neural information processing. In this first-ever book on complex-valued neural networks, the most active scientists at the forefront of the field describe theories and applications from various points of view to provide academic and industrial researchers with a comprehensive understanding of the fundamentals, features and prospects of the powerful complex-valued networks.

Data Stream Mining & Processing

Cyber Security Cryptography and Machine Learning

Artificial Intelligence and Machine Learning for EDGE Computing

Reinforcement Learning, second edition

Machine Learning for Financial Engineering

Handbook Of Machine Learning - Volume 1: Foundation Of Artif

20th International Conference, Manchester, UK, November 14–16, 2019, Proceedings, Part I

Artificial Intelligence and Machine Learning for Predictive and Analytical Rendering in Edge Computing focuses on the role of AI and machine learning as it impacts and works alongside Edge Computing. Sections cover the growing number of devices and applications in diversified domains of industry, including gaming, speech recognition, medical diagnostics, robotics and computer vision and how they are being driven by Big Data, Artificial Intelligence, Machine Learning and distributed computing, may it be Cloud Computing or the evolving Fog and Edge Computing paradigms. Challenges covered include remote storage and computing, bandwidth overload due to transportation of data from End nodes to Cloud leading in latency issues, security issues in transporting sensitive medical and financial information across larger gaps in points of data generation and computing, as well as design features of Edge nodes to store and run AI/ML algorithms for effective rendering. Provides a reference handbook on the evolution of distributed systems, including Cloud, Fog and Edge Computing Integrates the various Artificial Intelligence and Machine Learning techniques for effective predictions at Edge rather than Cloud or remote Data Centers Provides insight into the features and constraints in Edge Computing and storage, including hardware constraints and the technological/architectural developments that shall overcome those constraints

This book presents the proceedings of the 24th European Conference on Artificial Intelligence (ECAI 2020), held in Santiago de Compostela, Spain, from 29 August to 8 September 2020. The conference was postponed from June, and much of it conducted online due to the COVID-19 restrictions. The conference is one of the principal occasions for researchers and practitioners of AI to meet and discuss the latest trends and challenges in all fields of AI and to demonstrate innovative applications and uses of advanced AI technology. The book also includes the proceedings of the 10th Conference on Prestigious Applications of Artificial Intelligence (PAIS 2020) held at the same time. A record number of more than 1,700 submissions was received for ECAI 2020, of which 1,443 were reviewed. Of these, 361 full-papers and 36 highlight papers were accepted (an acceptance rate of 25% for full-papers and 45% for highlight papers). The book is divided into three sections: ECAI full papers; ECAI highlight papers; and PAIS papers. The topics of these papers cover all aspects of AI, including Agent-based and Multi-agent Systems; Computational Intelligence; Constraints and Satisfiability; Games and Virtual Environments; Heuristic Search; Human Aspects in AI; Information Retrieval and Filtering; Knowledge Representation and Reasoning; Machine Learning; Multidisciplinary Topics and Applications; Natural Language Processing; Planning and Scheduling; Robotics; Safe, Explainable, and Trustworthy AI; Semantic Technologies; Uncertainty in AI; and Vision. The book will be of interest to all those whose work involves the use of AI technology.

Recent decades have seen rapid advances in automatization processes, supported by modern machines and computers. The result is significant increases in system complexity and state changes, information sources, the need for faster data handling and the integration of environmental influences. Intelligent systems, equipped with a taxonomy of data-driven system identification and machine learning algorithms, can handle these problems partially. Conventional learning algorithms in a batch off-line setting fail whenever dynamic changes of the process appear due to non-stationary environments and external influences. Learning in Non-Stationary Environments: Methods and Applications offers a wide-ranging, comprehensive review of recent developments and important methodologies in the field. The coverage focuses on dynamic learning in unsupervised problems, dynamic learning in supervised classification and dynamic learning in supervised regression problems. A later section is dedicated to applications in which dynamic learning methods serve as keystones for achieving models with high accuracy. Rather than rely on a mathematical theorem/proof style, the editors highlight numerous figures, tables, examples and applications, together with their explanations. This approach offers a useful basis for further investigation and fresh ideas and motivates and inspires newcomers to explore this promising and still emerging field of research.

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.

Statistical Machine Learning for Human Behaviour Analysis

Applications of Machine Learning

30th International Conference on Artificial Neural Networks, Bratislava, Slovakia, September 14–17, 2021, Proceedings, Part II

Safeguard your system by making your machines intelligent using the Python ecosystem

An Introduction To Machine Learning In Quantitative Finance

Machine Learning and Knowledge Discovery in Databases

Deep Learning and Neural Networks

Forecasting is required in many situations. Stocking an inventory may require forecasts of demand months in advance. Telecommunication routing requires traffic forecasts a few minutes ahead. Whatever the circumstances or time horizons involved, forecasting is an important aid in effective and efficient planning. This textbook provides a comprehensive introduction to forecasting methods and presents enough information about each method for readers to use them sensibly.

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.

This book covers applications of machine learning in artificial intelligence. The specific topics covered include human language, heterogeneous and streaming data, unmanned systems, neural information processing, marketing and the social sciences, bioinformatics and robotics, etc. It also provides a broad range of techniques that can be successfully applied and adopted in different areas. Accordingly, the book offers an interesting and insightful read for scholars in the areas of computer vision, speech recognition, healthcare, business, marketing, and bioinformatics.

Deep learning methods offer a lot of promise for time series forecasting, such as the automatic learning of temporal dependence and the automatic handling of temporal structures like trends and seasonality. With clear explanations, standard Python libraries, and step-by-step tutorial lessons you'll discover how to develop deep learning models for your own time series forecasting projects.

Gaussian Processes for Machine Learning

Learning and Decision-Making from Rank Data

A Comparison of Machine Learning Model Validation Schemes for Non-stationary Time Series Data

ECAI 2020

European Conference, ECML PKDD 2016, Riva del Garda, Italy, September 19-23, 2016, Proceedings, Part III

9789386173423

Statistical Reinforcement Learning

This two-volume set of LNCS 11871 and 11872 constitutes the thoroughly refereed conference proceedings of the 20th International Conference on Intelligent Data Engineering and Automated Learning, IDEAL 2019, held in Manchester, UK, in November 2019. The 94 full papers presented were carefully reviewed and selected from 149 submissions. These papers provided a timely sample of the latest advances in data engineering and machine learning, from methodologies, frameworks, and algorithms to applications. The core themes of IDEAL 2019 include big data challenges, machine learning, data mining, information retrieval and management, bio-/neuro-informatics, bio-inspired models (including neural networks, evolutionary computation and swarm intelligence), agents and hybrid intelligent systems, real-world applications of intelligent techniques and AI.

This book constitutes the proceedings of the third International Conference on Data Stream and Mining and Processing, DSMP 2020, held in Lviv, Ukraine*, in August 2020. The 36 full papers presented in this volume were carefully reviewed and selected from 134 submissions. The papers are organized in topical sections of hybrid systems of computational intelligence; machine vision and pattern recognition; dynamic data mining & data stream mining; big data & data science using intelligent approaches. *The conference was held virtually due to the COVID-19 pandemic.

This book constitutes the refereed proceedings of the joint conference on Machine Learning and Knowledge Discovery in Databases: ECML PKDD 2008, held in Antwerp, Belgium, in September 2008. The 100 papers presented in two volumes, together with 5 invited talks, were carefully reviewed and selected from 521 submissions. In addition to the regular papers the volume contains 14 abstracts of papers appearing in full version in the Machine Learning Journal and the Knowledge Discovery and Databases Journal of Springer. The conference intends to provide an international forum for the discussion of the latest high quality research results in all areas related to machine learning and knowledge discovery in databases. The topics addressed are application of machine learning and data mining methods to real-world problems, particularly exploratory research that describes novel learning and mining tasks and applications requiring non-standard techniques.

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.

Third International Conference, DSMP 2020, Lviv, Ukraine, August 21-25, 2020, Proceedings

Predict the Future with MLPs, CNNs and LSTMs in Python

Mathematics for Machine Learning

24th European Conference on Artificial Intelligence, 29 August-8 September 2020, Santiago de Compostela, Spain - Including 10th Conference on Prestigious Applications of Artificial Intelligence (PAIS 2020)

Modern Machine Learning Approaches

Proceedings of the Eighth International Conference

From Theory to Practice

This book constitutes selected papers from the First International Workshop on Deployable Machine Learning for Security Defense, MLHat 2020, held in August 2020. Due to the COVID-19 pandemic the conference was held online. The 8 full papers were thoroughly reviewed and selected from 13 qualified submissions. The papers are organized in the following topical sections: understanding the adversaries; adversarial ML for better security; threats on networks.

A comprehensive introduction to machine learning that uses probabilistic models and inference as a unifying approach. Today's Web-enabled deluge of electronic data calls for automated methods of data analysis. Machine learning provides these, developing methods that can automatically detect patterns in data and then use the uncovered patterns to predict future data. This textbook offers a comprehensive and self-contained introduction to the field of machine learning, based on a unified, probabilistic approach. The coverage combines breadth and depth, offering necessary background material on such topics as probability, optimization, and linear algebra as well as discussion of recent developments in the field, including conditional random fields, L1 regularization, and deep learning. The book is written in an informal, accessible style, complete with pseudo-code for the most important algorithms. All topics are copiously illustrated with color images and worked examples drawn from such application domains as biology, text processing, computer vision, and robotics. Rather than providing a cookbook of different heuristic methods, the book stresses a principled model-based approach, often using the language of graphical models to specify models in a concise and intuitive way. Almost all the models described have been implemented in a MATLAB software package—PMTK (probabilistic modeling toolkit)—that is freely available online. The book is suitable for upper-level undergraduates with an introductory-level college math background and beginning graduate students.

Get into the world of smart data security using machine learning algorithms and Python libraries Key Features Learn machine learning algorithms and cybersecurity fundamentals Automate your daily workflow by applying use cases to many facets of security Implement smart machine learning solutions to detect various cybersecurity problems Book Description Cyber threats today are one of the costliest losses that an organization can face. In this book, we use the most efficient tool to solve the big problems that exist in the cybersecurity domain. The book begins by giving you the basics of ML in cybersecurity using Python and its libraries. You will explore various ML domains (such as time series analysis and ensemble modeling) to get your foundations right. You will implement various examples such as building system to identify malicious URLs, and building a program to detect fraudulent emails and spam. Later, you will learn how to make effective use of K-means algorithm to develop a solution to detect and alert you to any malicious activity in the network. Also learn how to implement biometrics and fingerprint to validate whether the user is a legitimate user or not. Finally, you will see how we change the game with TensorFlow and learn how deep learning is effective for creating models and training systems What you will learn Use machine learning algorithms with complex datasets to implement cybersecurity concepts Implement machine learning algorithms such as clustering, k-means, and Naive Bayes to solve real-world problems Learn to speed up a system using Python libraries with NumPy, Scikit-learn, and CUDA Understand how to combat malware, detect spam, and fight financial fraud to mitigate cyber crimes Use TensorFlow in the cybersecurity domain and implement real-world examples Learn how machine learning and Python can be used in complex cyber issues Who this book is for This book is for the data scientists, machine learning developers, security researchers, and anyone keen to apply machine learning to up-skill computer security. Having some working knowledge of Python and being familiar with the basics of machine learning and cybersecurity fundamentals will help to get the most out of the book

The three volume set LNAI 9851, LNAI 9852, and LNAI 9853 constitutes the refereed proceedings of the European Conference on Machine Learning and Knowledge Discovery in Databases, ECML PKDD 2016, held in Riva del Garda, Italy, in September 2016. The 123 full papers and 16 short papers presented were carefully reviewed and selected from a total of 460 submissions. The papers presented focus on practical and real-world studies of machine learning, knowledge discovery, data mining; innovative prototype implementations or mature systems that use machine learning techniques and knowledge discovery processes in a real setting; recent advances at the frontier of machine learning and data mining with other disciplines. Part I and Part II of the proceedings contain the full papers of the contributions presented in the scientific track and abstracts of the scientific plenary talks. Part III contains the full papers of the contributions presented in the industrial track, short papers describing demonstration, the nectar papers, and the abstracts of the industrial plenary talks.

5th International Symposium, CSCML 2021, Be'er Sheva, Israel, July 8–9, 2021, Proceedings

First International Workshop, MLHat 2020, San Diego, CA, USA, August 24, 2020, Proceedings

A Unified Framework

Metaheuristics

Control Charts and Machine Learning for Anomaly Detection in Manufacturing

21st International Conference on Artificial Neural Networks, Espoo, Finland, June 14-17, 2011, Proceedings, Part II

Machine Learning in Non-stationary Environments

Machine Learning Proceedings 1994

With the wide availability of large amounts of data and acute need for extracting useful information from such data, intelligent data analysis has attracted great attention and contributed to solving many practical tasks, ranging from scientific research, industrial process and daily life. In many cases the data evolve over time or change from one domain to another. The non-stationary nature of the data brings a new challenge for many existing learning algorithms, which are based on the stationary assumption. This dissertation addresses three crucial problems towards the effective handling of non-stationary data by investigating systematic methods for sample reweighting. Sample reweighting is a problem that infers sample-dependent weights for a data collection to match another data collection which exhibits distributional difference. It is known as the density-ratio estimation problem and the estimation results can be used in several machine learning tasks. This research proposes a set of methods for distribution matching by developing novel density-ratio methods that incorporate the characters of different non-stationary data analysis tasks. The contributions are summarized below. First, for the domain adaptation of classification problems a novel discriminative density-ratio method is proposed. This approach combines three learning objectives: minimizing generalized risk on the reweighted training data, minimizing class-wise distribution discrepancy and maximizing the separation margin on the test data. To solve the discriminative density-ratio problem, two algorithms are presented on the basis of a block coordinate update optimization scheme. Experiments conducted on different domain adaptation scenarios demonstrate the effectiveness of the proposed algorithms. Second, for detecting novel instances in the test data a locally-adaptive kernel density-ratio method is proposed. While traditional novelty detection algorithms are limited to detect either emerging novel instances which are completely new, or evolving novel instances whose distribution are different from previously-seen ones, the proposed algorithm builds on the success of the idea of using density ratio as a measure of evolving novelty and augments with structural information of each data instance's neighborhood. This makes the estimation of density ratio more reliable, and results in detection of emerging as well as evolving novelties. In addition, the proposed locally-adaptive kernel novelty detection method is applied in the social media analysis and shows favorable performance over other existing approaches. As the time continuity of social media streams, the novelty is usually characterized by the combination of emerging and evolving. One reason is the existence of large common vocabularies between different topics. Another reason is that there are high possibilities of topics being continuously discussed in sequential batch of collections, but showing different level of intensity. Thus, the presented novelty detection algorithm demonstrates its effectiveness in the social media data analysis. Lastly, an auto-tuning method for the non-parametric kernel mean matching estimator is presented. It introduces a new quality measure for evaluating the goodness of distribution matching which reflects the normalized mean square error of estimates. The proposed quality measure does not depend on the learner in the following step and accordingly allows the model selection procedures for importance estimation and prediction model learning to be completely separated.

Combinatorial optimization is the process of finding the best, or optimal, so lution for problems with a discrete set of feasible solutions. Applications arise in numerous settings involving operations management and logistics, such as routing, scheduling, packing, inventory and production management, lo cation, logic, and assignment of resources. The economic impact of combi natorial optimization is profound, affecting sectors as diverse as transporta tion (airlines, trucking, rail, and shipping), forestry, manufacturing, logistics, aerospace, energy (electrical power, petroleum, and natural gas), telecommu nications, biotechnology, financial services, and agriculture. While much progress has been made in finding exact (provably optimal) so lutions to some combinatorial optimization problems, using techniques such as dynamic programming, cutting planes, and branch and cut methods, many hard combinatorial problems are still not solved exactly and require good heuristic methods. Moreover, reaching "optimal solutions" is in many cases meaningless, as in practice we are often dealing with models that are rough simplifications of reality. The aim of heuristic methods for combinatorial op timization is to quickly produce good-quality solutions, without necessarily providing any guarantee of solution quality. Metaheuristics are high level procedures that coordinate simple heuristics, such as local search, to find solu tions that are of better quality than those found by the simple heuristics alone: Modem metaheuristics include simulated annealing, genetic algorithms, tabu search, GRASP, scatter search, ant colony optimization, variable neighborhood search, and their hybrids.

This book introduces the latest research on advanced control charts and new machine learning approaches to detect abnormalities in the smart manufacturing process. By approaching anomaly detection using both statistics and machine learning, the book promotes interdisciplinary cooperation between the research communities, to jointly develop new anomaly detection approaches that are more suitable for the 4.0 Industrial Revolution. The book provides ready-to-use algorithms and parameter sheets, enabling readers to design advanced control charts and machine learning-based approaches for anomaly detection in manufacturing. Case studies are introduced in each chapter to help practitioners easily apply these tools to real-world manufacturing processes. The book is of interest to researchers, industrial experts, and postgraduate students in the fields of industrial engineering, automation, statistical learning, and manufacturing industries.

Machine Learning in Non-Stationary Environments

A Probabilistic Perspective

Machine Learning

Deep Learning for Time Series Forecasting

Hands-On Machine Learning for Cybersecurity

Adaptive Learning Algorithms for Non-stationary Data

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.

Theory, algorithms, and applications of machine learning techniques to overcome "covariate shift" non-stationarity.

The two-volume set LNCS 7552 + 7553 constitutes the proceedings of the 22nd International Conference on Artificial Neural Networks, ICANN 2012, held in Lausanne, Switzerland, in September 2012. The 162 papers included in the proceedings were carefully reviewed and selected from 247 submissions. They are organized in topical sections named: theoretical neural computation; information and optimization; from neurons to neuromorphism; spiking dynamics; from single neurons to networks; complex firing patterns; movement and motion; from sensation to perception; object and face recognition; reinforcement learning; bayesian and echo state networks; recurrent neural networks and reservoir computing; coding architectures; interacting with the brain; swarm intelligence and decision-making; mulitlayer perceptrons and kernel networks; training and learning; inference and recognition; support vector machines; self-organizing maps and clustering; clustering, mining and exploratory analysis; bioinformatics; and time weries and forecasting.