Bayesian Deep Learning Uncertainty In Deep Learning

Yeah, reviewing a ebook **Bayesian Deep Learning Uncertainty In Deep Learning** could increase your near friends listings. This is just one of the solutions for you to be successful. As understood, finishing does not suggest that you have wonderful points.

Comprehending as with ease as arrangement even more than other will allow each success. adjacent to, the statement as capably as perception of this Bayesian Deep Learning Uncertainty In Deep Learning can be taken as well as picked to act.

Advanced Lectures on Machine Learning - Olivier Bousquet 2011-03-22
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.

Bayesian Deep Learning and Uncertainty in Computer Vision - Buu Truong Phan 2019
Visual data contains rich information about the operating environment of an intelligent robotic system. Extracting this information allows intelligent systems to reason and decide their future actions. Erroneous visual information, therefore, can lead to poor decisions, causing accidents and casualties, especially in a safety-critical application such as automated driving. One way to prevent this is by measuring the level of uncertainty in the visual information interpretation, so that the system knows the reliability degree of the extracted information. Deep neural networks are now being used in many vision tasks due to their superior accuracy compared to traditional machine learning methods. However, their estimated uncertainties have been shown to be unreliable. To mitigate this issue, researchers have developed methods and tools to apply Bayesian modeling to deep neural networks. This results in a class of models known as Bayesian neural networks, whose uncertainty estimates are more reliable and informative. In this thesis, we make the following contributions in the context of Bayesian Neural Network applied to vision tasks. In particular: - We improve the understanding of visual uncertainty estimates from Bayesian deep models. Specifically, we study

the behavior of Bayesian deep models applied to road-scene image segmentation under different factors, such as varying weather, depth, and occlusion levels. - We show the importance of model calibration technique in the context of autonomous driving, which strengthens the reliability of the estimated uncertainty. We demonstrate its effectiveness in a simple object localization task. - We address the high run-time cost of the current Bayesian deep learning techniques. We develop a distillation technique based on the Dirichlet distribution, which allows us to estimate the uncertainties in real-time.

Probabilistic Machine Learning - Kevin P. Murphy 2023-08-15

Probabilistic Modelling of Uncertainty with Bayesian Nonparametric Machine Learning - Charles Gadd 2018

Probabilistic Deep Learning - Oliver Duerr 2020-11-10

Probabilistic Deep Learning is a hands-on guide to the principles that support neural networks. Learn to improve network performance with the right distribution for different data types, and discover Bayesian variants that can state their own uncertainty to increase accuracy. This book provides easy-to-apply code and uses popular frameworks to keep you focused on practical applications. Summary Probabilistic Deep Learning: With Python, Keras and TensorFlow Probability teaches the increasingly popular probabilistic approach to deep learning that allows you to refine your results more quickly and accurately without much trial-and-error testing. Emphasizing practical techniques that use the Python-based Tensorflow Probability Framework, you'll learn to build highly-performant deep learning applications that can reliably handle the noise and uncertainty of real-world data. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the technology The world is a noisy and uncertain place. Probabilistic deep learning models capture that noise and uncertainty, pulling it into real-world scenarios. Crucial for self-driving cars and scientific testing, these techniques help deep learning engineers assess the accuracy of their results, spot errors, and improve their understanding of how algorithms work. About the book Probabilistic Deep Learning is a hands-on guide to the principles that support neural networks. Learn to improve network performance with the right distribution for different data types, and discover Bayesian variants that can state their own uncertainty to increase accuracy. This book provides easy-to-apply code and uses popular frameworks to keep you focused on practical applications. What's inside Explore maximum likelihood and the statistical basis of deep learning Discover probabilistic models that can indicate possible outcomes Learn to use normalizing flows for modeling and generating complex distributions Use Bayesian neural networks to access the uncertainty in the model About the reader For experienced machine learning developers. About the author Oliver Dürr is a professor at the University of Applied Sciences in Konstanz, Germany. Beate Sick holds a chair for applied statistics at ZHAW and works as a researcher and lecturer at the University of Zurich. Elvis Murina is a data scientist. Table of Contents PART 1 - BASICS OF DEEP LEARNING 1 Introduction to probabilistic deep learning 2 Neural network architectures 3 Principles of curve fitting PART 2 - MAXIMUM

LIKELIHOOD APPROACHES FOR PROBABILISTIC DL MODELS 4 Building loss functions with the likelihood approach 5 Probabilistic deep learning models with TensorFlow Probability 6 Probabilistic deep learning models in the wild PART 3 - BAYESIAN APPROACHES FOR PROBABILISTIC DL MODELS 7 Bayesian learning 8 Bayesian neural networks

Uncertainty for Safe Utilization of Machine Learning in Medical Imaging, and Graphs in Biomedical Image Analysis - Carole H. Sudre 2020-10-05

This book constitutes the refereed proceedings of the Second International Workshop on Uncertainty for Safe Utilization of Machine Learning in Medical Imaging, UNSURE 2020, and the Third International Workshop on Graphs in Biomedical Image Analysis, GRAIL 2020, held in conjunction with MICCAI 2020, in Lima, Peru, in October 2020. The workshops were held virtually due to the COVID-19 pandemic. For UNSURE 2020, 10 papers from 18 submissions were accepted for publication. They focus on developing awareness and encouraging research in the field of uncertainty modelling to enable safe implementation of machine learning tools in the clinical world. GRAIL 2020 accepted 10 papers from the 12 submissions received. The workshop aims to bring together scientists that use and develop graph-based models for the analysis of biomedical images and to encourage the exploration of graph-based models for difficult clinical problems within a variety of biomedical imaging contexts.

The Alignment Problem: Machine Learning and Human Values - Brian Christian 2020-10-06

A jaw-dropping exploration of everything that goes wrong when we build AI systems and the movement to fix them. Today’s “machine-learning” systems, trained by data, are so effective that we’ve invited them to see and hear for us—and to make decisions on our behalf. But alarm bells are ringing. Recent years have seen an eruption of concern as the field of machine learning advances. When the systems we attempt to teach will not, in the end, do what we want or what we expect, ethical and potentially existential risks emerge. Researchers call this the alignment problem. Systems cull résumés until, years later, we discover that they have inherent gender biases. Algorithms decide bail and parole—and appear to assess Black and White defendants differently. We can no longer assume that our mortgage application, or even our medical tests, will be seen by human eyes. And as autonomous vehicles share our streets, we are increasingly putting our lives in their hands. The mathematical and computational models driving these changes range in complexity from something that can fit on a spreadsheet to a complex system that might credibly be called “artificial intelligence.” They are steadily replacing both human judgment and explicitly programmed software. In best-selling author Brian Christian’s riveting account, we meet the alignment problem’s “first-responders,” and learn their ambitious plan to solve it before our hands are completely off the wheel. In a masterful blend of history and on-the ground reporting, Christian traces the explosive growth in the field of machine learning and surveys its current, sprawling frontier. Readers encounter a discipline finding its legs amid exhilarating and sometimes terrifying progress. Whether they—and we—succeed or fail in solving the alignment problem will be a defining human story. The Alignment Problem offers an unflinching reckoning with humanity’s biases and blind spots, our own

unstated assumptions and often contradictory goals. A dazzlingly interdisciplinary work, it takes a hard look not only at our technology but at our culture—and finds a story by turns harrowing and hopeful.

Knowledge Guided Machine Learning - Anuj Karpatne 2022-08-15

Given their tremendous success in commercial applications, machine learning (ML) models are increasingly being considered as alternatives to science-based models in many disciplines. Yet, these "black-box" ML models have found limited success due to their inability to work well in the presence of limited training data and generalize to unseen scenarios. As a result, there is a growing interest in the scientific community on creating a new generation of methods that integrate scientific knowledge in ML frameworks. This emerging field, called scientific knowledge-guided ML (KGML), seeks a distinct departure from existing "data-only" or "scientific knowledge-only" methods to use knowledge and data at an equal footing. Indeed, KGML involves diverse scientific and ML communities, where researchers and practitioners from various backgrounds and application domains are continually adding richness to the problem formulations and research methods in this emerging field. Knowledge Guided Machine Learning: Accelerating Discovery using Scientific Knowledge and Data provides an introduction to this rapidly growing field by discussing some of the common themes of research in KGML using illustrative examples, case studies, and reviews from diverse application domains and research communities as book chapters by leading researchers. KEY FEATURES First-of-its-kind book in an emerging area of research that is gaining widespread attention in the scientific and data science fields Accessible to a broad audience in data science and scientific and engineering fields Provides a coherent organizational structure to the problem formulations and research methods in the emerging field of KGML using illustrative examples from diverse application domains Contains chapters by leading researchers, which illustrate the cutting-edge research trends, opportunities, and challenges in KGML research from multiple perspectives Enables cross-pollination of KGML problem formulations and research methods across disciplines Highlights critical gaps that require further investigation by the broader community of researchers and practitioners to realize the full potential of KGML

Variational Methods for Machine Learning with Applications to Deep Networks - Lucas Pinheiro Cinelli 2021-05-10

This book provides a straightforward look at the concepts, algorithms and advantages of Bayesian Deep Learning and Deep Generative Models. Starting from the model-based approach to Machine Learning, the authors motivate Probabilistic Graphical Models and show how Bayesian inference naturally lends itself to this framework. The authors present detailed explanations of the main modern algorithms on variational approximations for Bayesian inference in neural networks. Each algorithm of this selected set develops a distinct aspect of the theory. The book builds from the ground-up well-known deep generative models, such as Variational Autoencoder and subsequent theoretical developments. By also exposing the main issues of the algorithms together with different methods to mitigate such issues, the book supplies the necessary knowledge on generative

models for the reader to handle a wide range of data types: sequential or not, continuous or not, labelled or not. The book is self-contained, promptly covering all necessary theory so that the reader does not have to search for additional information elsewhere. Offers a concise self-contained resource, covering the basic concepts to the algorithms for Bayesian Deep Learning; Presents Statistical Inference concepts, offering a set of elucidative examples, practical aspects, and pseudo-codes; Every chapter includes hands-on examples and exercises and a website features lecture slides, additional examples, and other support material.

Graphical Models, Exponential Families, and Variational Inference - Martin J. Wainwright 2008

The core of this paper is a general set of variational principles for the problems of computing marginal probabilities and modes, applicable to multivariate statistical models in the exponential family.

Pattern Recognition and Machine Learning - Christopher M. Bishop 2016-08-23

This is the first textbook on pattern recognition to present the Bayesian viewpoint. The book presents approximate inference algorithms that permit fast approximate answers in situations where exact answers are not feasible. It uses graphical models to describe probability distributions when no other books apply graphical models to machine learning. No previous knowledge of pattern recognition or machine learning concepts is assumed. Familiarity with multivariate calculus and basic linear algebra is required, and some experience in the use of probabilities would be helpful though not essential as the book includes a self-contained introduction to basic probability theory.

Uncertainty Estimation for Dense Stereo Matching Using Bayesian Deep Learning - Max Mehltretter 2021

Fully Bayesian Learning and Classic Deep Learning - Elio Abi Younes 2020

"Classic deep learning algorithms are powerful tools for the construction of accurate predictive models for labeled data. However, traditional deep neural networks designed to learning such models are both prone to overfitting and incapable of assessing uncertainty. In contrast, Bayesian learning based upon the emergence of Markov chain Monte Carlo methods and variational inference provides strong ability to express uncertainty in predictions and improve the estimated posterior probability based on new evidence. This work further assesses the efficiency and accuracy of Bayesian inference in complex settings. We provide an in-depth empirical analysis of the methods on both real and synthetic data in the context of regression and image classification. Specifically, we develop a unified Bayesian deep neural network model interleaving Bayesian sampling into deep learning. By rephrasing these learning techniques upon a common theoretical ground casting (1) the application of fully Bayesian learning for deep neural networks rather than pure optimization-based or approximate learning and (2) the most significant regularization technique in neural networks, dropout, as approximate Bayesian inference, we perform a clear comparison proving the efficiency of Bayesian deep learning to maintain state-of-the-art performance compared to existing methods while mitigating the problem of uncertainty in deep learning"--

A Tutorial on Thompson Sampling - Daniel J. Russo 2018-07-12

Thompson sampling is an algorithm for online decision problems where actions are taken sequentially in a manner that must balance between exploiting what is known to maximize immediate performance and investing to accumulate new information that may improve future performance. The algorithm addresses a broad range of problems in a computationally efficient manner and is therefore enjoying wide use. A Tutorial on Thompson Sampling covers the algorithm and its application, illustrating concepts through a range of examples, including Bernoulli bandit problems, shortest path problems, product recommendation, assortment, active learning with neural networks, and reinforcement learning in Markov decision processes. Most of these problems involve complex information structures, where information revealed by taking an action informs beliefs about other actions. It also discusses when and why Thompson sampling is or is not effective and relations to alternative algorithms.

Physics-aware, Bayesian Machine Learning Models for Uncertainty Quantification of High-dimensional Systems in the Small Data Regime - Constantin Georg Grigo 2020

Machine Learning - Kevin P. Murphy 2012-08-24

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.

Information Processing and Management of Uncertainty in Knowledge-Based Systems - Marie-Jeanne Lesot 2020-06-05

This three volume set (CCIS 1237-1239) constitutes the proceedings of the 18th International Conference on Information Processing and Management of Uncertainty in Knowledge-Based Systems, IPMU 2020, in June 2020. The conference was scheduled to take place in Lisbon, Portugal, at University of Lisbon, but due to COVID-19 pandemic it was held virtually. The 173 papers were carefully reviewed and selected from 213 submissions. The papers are organized in topical sections: homage

to Enrique Ruspini; invited talks; foundations and mathematics; decision making, preferences and votes; optimization and uncertainty; games; real world applications; knowledge processing and creation; machine learning I; machine learning II; XAI; image processing; temporal data processing; text analysis and processing; fuzzy interval analysis; theoretical and applied aspects of imprecise probabilities; similarities in artificial intelligence; belief function theory and its applications; aggregation: theory and practice; aggregation: pre-aggregation functions and other generalizations of monotonicity; aggregation: aggregation of different data structures; fuzzy methods in data mining and knowledge discovery; computational intelligence for logistics and transportation problems; fuzzy implication functions; soft methods in statistics and data analysis; image understanding and explainable AI; fuzzy and generalized quantifier theory; mathematical methods towards dealing with uncertainty in applied sciences; statistical image processing and analysis, with applications in neuroimaging; interval uncertainty; discrete models and computational intelligence; current techniques to model, process and describe time series; mathematical fuzzy logic and graded reasoning models; formal concept analysis, rough sets, general operators and related topics; computational intelligence methods in information modelling, representation and processing.

Bayesian Nonparametrics via Neural Networks - Herbert K. H. Lee 2004-01-01

Bayesian Nonparametrics via Neural Networks is the first book to focus on neural networks in the context of nonparametric regression and classification, working within the Bayesian paradigm. Its goal is to demystify neural networks, putting them firmly in a statistical context rather than treating them as a black box. This approach is in contrast to existing books, which tend to treat neural networks as a machine learning algorithm instead of a statistical model. Once this underlying statistical model is recognized, other standard statistical techniques can be applied to improve the model. The Bayesian approach allows better accounting for uncertainty. This book covers uncertainty in model choice and methods to deal with this issue, exploring a number of ideas from statistics and machine learning. A detailed discussion on the choice of prior and new noninformative priors is included, along with a substantial literature review. Written for statisticians using statistical terminology, Bayesian Nonparametrics via Neural Networks will lead statisticians to an increased understanding of the neural network model and its applicability to real-world problems.

Towards Intelligent Operation of Future Power System - Tingqi Zhang 2022

Mathematics for Machine Learning - Marc Peter Deisenroth 2020-04-23

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.

Advances in Bayesian Machine Learning - Chao Ma 2021

Uncertainty for Safe Utilization of Machine Learning in Medical Imaging and Clinical Image-Based Procedures - Hayit Greenspan 2019-10-10

This book constitutes the refereed proceedings of the First International Workshop on Uncertainty for Safe Utilization of Machine Learning in Medical Imaging, UNSURE 2019, and the 8th International Workshop on Clinical Image-Based Procedures, CLIP 2019, held in conjunction with MICCAI 2019, in Shenzhen, China, in October 2019. For UNSURE 2019, 8 papers from 15 submissions were accepted for publication. They focus on developing awareness and encouraging research in the field of uncertainty modelling to enable safe implementation of machine learning tools in the clinical world. CLIP 2019 accepted 11 papers from the 15 submissions received. The workshops provides a forum for work centred on specific clinical applications, including techniques and procedures based on comprehensive clinical image and other data.

Computer Safety, Reliability, and Security - Alexander Romanovsky 2019-09-02

This book constitutes the proceedings of the Workshops held in conjunction with SAFECOMP 2019, 38th International Conference on Computer Safety, Reliability and Security, in September 2019 in Turku, Finland. The 32 regular papers included in this volume were carefully reviewed and selected from 43 submissions; the book also contains two invited papers. The workshops included in this volume are: ASSURE 2019: 7th International Workshop on Assurance Cases for Software-Intensive Systems DECSoS 2019: 14th ERCIM/EWICS/ARTEMIS Workshop on Dependable Smart Embedded and Cyber-Physical Systems and Systems-of-Systems SASSUR 2019: 8th International Workshop on Next Generation of System Assurance Approaches for Safety-Critical Systems STRIVE 2019: Second International Workshop on Safety, securiTy, and pRivacy In automotiVe systEms WAISE 2019: Second International Workshop on Artificial Intelligence Safety Engineering

Deep Learning for Unmanned Systems - Anis Koubaa 2021-10-01

This book is used at the graduate or advanced undergraduate level and many others. Manned and unmanned ground, aerial and marine vehicles enable many promising and revolutionary civilian and military applications that will change our life in the near future. These applications include, but are not limited to, surveillance, search and rescue, environment monitoring, infrastructure monitoring, self-driving cars, contactless last-mile delivery vehicles, autonomous ships, precision agriculture and transmission line inspection to name just a few. These vehicles will benefit from advances of deep learning as a subfield of machine learning able to endow these vehicles with different capability such as perception, situation

awareness, planning and intelligent control. Deep learning models also have the ability to generate actionable insights into the complex structures of large data sets. In recent years, deep learning research has received an increasing amount of attention from researchers in academia, government laboratories and industry. These research activities have borne some fruit in tackling some of the challenging problems of manned and unmanned ground, aerial and marine vehicles that are still open. Moreover, deep learning methods have been recently actively developed in other areas of machine learning, including reinforcement training and transfer/meta-learning, whereas standard, deep learning methods such as recent neural network (RNN) and coevolutionary neural networks (CNN). The book is primarily meant for researchers from academia and industry, who are working on in the research areas such as engineering, control engineering, robotics, mechatronics, biomedical engineering, mechanical engineering and computer science. The book chapters deal with the recent research problems in the areas of reinforcement learning-based control of UAVs and deep learning for unmanned aerial systems (UAS) The book chapters present various techniques of deep learning for robotic applications. The book chapters contain a good literature survey with a long list of references. The book chapters are well written with a good exposition of the research problem, methodology, block diagrams and mathematical techniques. The book chapters are lucidly illustrated with numerical examples and simulations. The book chapters discuss details of applications and future research areas.

Uncertainty for Safe Utilization of Machine Learning in Medical Imaging - Carole H. Sudre 2022-09-17

This book constitutes the refereed proceedings of the Fourth Workshop on Uncertainty for Safe Utilization of Machine Learning in Medical Imaging, UNSURE 2022, held in conjunction with MICCAI 2022. The conference was hybrid event held from Singapore. For this workshop, 13 papers from 22 submissions were accepted for publication. They focus on developing awareness and encouraging research in the field of uncertainty modelling to enable safe implementation of machine learning tools in the clinical world.

Uncertainty Modeling for Data Mining - Zengchang Qin 2014-10-30

Machine learning and data mining are inseparably connected with uncertainty. The observable data for learning is usually imprecise, incomplete or noisy. Uncertainty Modeling for Data Mining: A Label Semantics Approach introduces 'label semantics', a fuzzy-logic-based theory for modeling uncertainty. Several new data mining algorithms based on label semantics are proposed and tested on real-world datasets. A prototype interpretation of label semantics and new prototype-based data mining algorithms are also discussed. This book offers a valuable resource for postgraduates, researchers and other professionals in the fields of data mining, fuzzy computing and uncertainty reasoning. Zengchang Qin is an associate professor at the School of Automation Science and Electrical Engineering, Beihang University, China; Yongchuan Tang is an associate professor at the College of Computer Science, Zhejiang University, China.

Encyclopedia of Machine Learning - Claude Sammut 2011-03-28

This comprehensive encyclopedia, in A-Z format, provides easy access to relevant information for those seeking entry into any aspect within the broad field of Machine Learning.

Most of the entries in this preeminent work include useful literature references.

Bayesian Reasoning and Machine Learning - David Barber 2012-02-02

A practical introduction perfect for final-year undergraduate and graduate students without a solid background in linear algebra and calculus.

Large-scale Kernel Machines - Léon Bottou 2007

Solutions for learning from large scale datasets, including kernel learning algorithms that scale linearly with the volume of the data and experiments carried out on realistically large datasets. Pervasive and networked computers have dramatically reduced the cost of collecting and distributing large datasets. In this context, machine learning algorithms that scale poorly could simply become irrelevant. We need learning algorithms that scale linearly with the volume of the data while maintaining enough statistical efficiency to outperform algorithms that simply process a random subset of the data. This volume offers researchers and engineers practical solutions for learning from large scale datasets, with detailed descriptions of algorithms and experiments carried out on realistically large datasets. At the same time it offers researchers information that can address the relative lack of theoretical grounding for many useful algorithms. After a detailed description of state-of-the-art support vector machine technology, an introduction of the essential concepts discussed in the volume, and a comparison of primal and dual optimization techniques, the book progresses from well-understood techniques to more novel and controversial approaches. Many contributors have made their code and data available online for further experimentation. Topics covered include fast implementations of known algorithms, approximations that are amenable to theoretical guarantees, and algorithms that perform well in practice but are difficult to analyze theoretically. Contributors Léon Bottou, Yoshua Bengio, Stéphane Canu, Eric Cosatto, Olivier Chapelle, Ronan Collobert, Dennis DeCoste, Ramani Duraiswami, Igor Durdanovic, Hans-Peter Graf, Arthur Gretton, Patrick Haffner, Stefanie Jegelka, Stephan Kanthak, S. Sathiya Keerthi, Yann LeCun, Chih-Jen Lin, Gaëlle Loosli, Joaquin Quiñonero-Candela, Carl Edward Rasmussen, Gunnar Rätsch, Vikas Chandrakant Raykar, Konrad Rieck, Vikas Sindhwani, Fabian Sinz, Sören Sonnenburg, Jason Weston, Christopher K. I. Williams, Elad Yom-Tov

Artificial Intelligence and Machine Learning Methods in COVID-19 and Related Health Diseases - Victor Chang 2022-06-28

This Springer book provides a perfect platform to submit chapters that discuss the prospective developments and innovative ideas in artificial intelligence and machine learning techniques in the diagnosis of COVID-19. COVID-19 is a huge challenge to humanity and the medical sciences. So far as of today, we have been unable to find a medical solution (Vaccine). However, globally, we are still managing the use of technology for our work, communications, analytics, and predictions with the use of advancement in data science, communication technologies (5G & Internet), and AI. Therefore, we might be able to continue and live safely with the use of research in advancements in data science, AI, machine learning, mobile apps, etc., until we can find a medical solution such as a vaccine. We have selected eleven chapters after the vigorous review process. Each chapter has

demonstrated the research contributions and research novelty. Each group of authors must fulfill strict requirements.

Bayesian Reinforcement Learning - Mohammad Ghavamzadeh 2015-11-18

Bayesian methods for machine learning have been widely investigated, yielding principled methods for incorporating prior information into inference algorithms. This monograph provides the reader with an in-depth review of the role of Bayesian methods for the reinforcement learning (RL) paradigm. The major incentives for incorporating Bayesian reasoning in RL are that it provides an elegant approach to action-selection (exploration/exploitation) as a function of the uncertainty in learning, and it provides a machinery to incorporate prior knowledge into the algorithms. Bayesian Reinforcement Learning: A Survey first discusses models and methods for Bayesian inference in the simple single-step Bandit model. It then reviews the extensive recent literature on Bayesian methods for model-based RL, where prior information can be expressed on the parameters of the Markov model. It also presents Bayesian methods for model-free RL, where priors are expressed over the value function or policy class. Bayesian Reinforcement Learning: A Survey is a comprehensive reference for students and researchers with an interest in Bayesian RL algorithms and their theoretical and empirical properties.

Artificial Intelligence For High Energy Physics - Paolo Calafiura 2022-01-05

The Higgs boson discovery at the Large Hadron Collider in 2012 relied on boosted decision trees. Since then, high energy physics (HEP) has applied modern machine learning (ML) techniques to all stages of the data analysis pipeline, from raw data processing to statistical analysis. The unique requirements of HEP data analysis, the availability of high-quality simulators, the complexity of the data structures (which rarely are image-like), the control of uncertainties expected from scientific measurements, and the exabyte-scale datasets require the development of HEP-specific ML techniques. While these developments proceed at full speed along many paths, the nineteen reviews in this book offer a self-contained, pedagogical introduction to ML models' real-life applications in HEP, written by some of the foremost experts in their area.

Teaching and Learning with Cases - Laurence E. Lynn Jr 1999-01-01

Lynn introduces readers to the case method of instruction popularized by the John F. Kennedy School of Government and the Harvard Business School. This is a practical, process-oriented guide to teaching, writing, and learning with the case method. Lynn integrates insight from literature with his own extensive experience as a case teacher and writer, and as a trainer of case teachers and case writers. Lynn selects the broadest possible context for discussing the use of cases in teaching for maximum appeal to instructors and learners in diverse fields.

Advances in Intelligent Data Analysis XVIII - Michael Berthold 2020-10-09

This open access book constitutes the proceedings of the 18th International Conference on Intelligent Data Analysis, IDA 2020, held in Konstanz, Germany, in April 2020. The 45 full papers presented in this volume were carefully reviewed and selected from 114 submissions. Advancing Intelligent Data Analysis requires novel, potentially game-changing ideas. IDA's mission is to promote ideas over performance: a solid

motivation can be as convincing as exhaustive empirical evaluation. This work was published by Saint Philip Street Press pursuant to a Creative Commons license permitting commercial use. All rights not granted by the work's license are retained by the author or authors.

Developing Deep Learning and Bayesian Deep Learning Based Models for MR Neuroimaging - Gengyan Zhao 2019

Magnetic resonance (MR) neuroimaging is an active field in investigating brain structures and functions. After decades of development, the whole pipeline of MR neuroimaging tends to become mature, but many essential steps still faces challenges and difficulties, especially in the accuracy of the image segmentation, image generation and data prediction. Recently, the revival of deep neural networks made immense progress in the field of machine learning. The proposal of Bayesian deep learning further enabled the ability of uncertainty generation in deep learning prediction. In this work, we proposed and developed different kinds of Bayesian neural networks to improve the accuracy of brain segmentation, brain image synthesis, and brain function related behavior prediction. To overcome the challenges in brain segmentation, we proposed a fully-automated brain extraction pipeline combining deep Bayesian convolutional neural network (CNN) and fully connected three-dimensional (3D) conditional random field (CRF). To increase the image synthesis accuracy and improve the calibration of the model uncertainty, we proposed a Bayesian conditional generative adversarial network (GAN). To improve the brain function related behavior prediction, we proposed a Bayesian deep neural network (DNN), and a feature extraction and ranking method for it. Experiments were done on real data to validate the proposed methods. The comparison between our methods and the state-of-the-arts showed that our methods can significantly improve the testing accuracy and the behavior of the model uncertainty generated by the Bayesian neural networks matches our expectation.

Machine Learning and Knowledge Discovery in Databases - Massih-Reza Amini 2023-03-17

The multi-volume set LNAI 13713 until 13718 constitutes the refereed proceedings of the European Conference on Machine Learning and Knowledge Discovery in Databases, ECML PKDD 2022, which took place in Grenoble, France, in September 2022. The 236 full papers presented in these proceedings were carefully reviewed and selected from a total of 1060 submissions. In addition, the proceedings include 17 Demo Track contributions. The volumes are organized in topical sections as follows: Part I: Clustering and dimensionality reduction; anomaly detection; interpretability and explainability; ranking and recommender systems; transfer and multitask learning; Part II: Networks and graphs; knowledge graphs; social network analysis; graph neural networks; natural language processing and text mining; conversational systems; Part III: Deep learning; robust and adversarial machine learning; generative models; computer vision; meta-learning, neural architecture search; Part IV: Reinforcement learning; multi-agent reinforcement learning; bandits and online learning; active and semi-supervised learning; private and federated learning; . Part V: Supervised learning; probabilistic inference; optimal transport; optimization; quantum, hardware; sustainability; Part VI: Time series; financial machine learning; applications; applications:

transportation; demo track.

Bayesian Learning for Neural Networks - Radford M. Neal 2012-12-06

Artificial "neural networks" are widely used as flexible models for classification and regression applications, but questions remain about how the power of these models can be safely exploited when training data is limited. This book demonstrates how Bayesian methods allow complex neural network models to be used without fear of the "overfitting" that can occur with traditional training methods. Insight into the nature of these complex Bayesian models is provided by a theoretical investigation of the priors over functions that underlie them. A practical implementation of Bayesian neural network learning using Markov chain Monte Carlo methods is also described, and software for it is freely available over the Internet. Presupposing only basic knowledge of probability and statistics, this book should be of interest to researchers in statistics, engineering, and artificial intelligence.

Information Theory, Inference and Learning Algorithms - David J. C. MacKay 2003-09-25

Information theory and inference, taught together in this exciting textbook, lie at the heart of many important areas of modern technology - communication, signal processing, data mining, machine learning, pattern recognition, computational neuroscience, bioinformatics and cryptography. The book introduces theory in tandem with applications. Information theory is taught alongside practical communication systems such as arithmetic coding for data compression and sparse-graph codes for error-correction. Inference techniques, including message-passing algorithms, Monte Carlo methods and variational approximations, are developed alongside applications to clustering, convolutional codes, independent component analysis, and neural networks. Uniquely, the book covers state-of-the-art error-correcting codes, including low-density-parity-check codes, turbo codes, and digital fountain codes - the twenty-first-century standards for satellite communications, disk drives, and data broadcast. Richly illustrated, filled with worked examples and over 400 exercises, some with detailed solutions, the book is ideal for self-learning, and for undergraduate or graduate courses. It also provides an unparalleled entry point for professionals in areas as diverse as computational biology, financial engineering and machine learning.

Reinforcement Learning - Marco Wiering 2012-03-05

Reinforcement learning encompasses both a science of adaptive behavior of rational beings in uncertain environments and a computational methodology for finding optimal behaviors for challenging problems in control, optimization and adaptive behavior of intelligent agents. As a field, reinforcement learning has progressed tremendously in the past decade. The main goal of this book is to present an up-to-date series of survey articles on the main contemporary sub-fields of reinforcement learning. This includes surveys on partially observable environments, hierarchical task decompositions, relational knowledge representation and predictive state representations. Furthermore, topics such as transfer, evolutionary methods and continuous spaces in reinforcement learning are surveyed. In addition, several chapters review reinforcement learning methods in robotics, in games, and in computational neuroscience. In total seventeen different subfields are

presented by mostly young experts in those areas, and together they truly represent a state-of-the-art of current reinforcement learning research. Marco Wiering works at the artificial intelligence department of the University of Groningen in the Netherlands. He has published extensively on various reinforcement learning topics. Martijn van Otterlo works in the cognitive artificial intelligence group at the Radboud University Nijmegen in The Netherlands. He has mainly focused on expressive knowledge representation in reinforcement learning settings.

Probability for Machine Learning - Jason Brownlee 2019-09-24
Probability is the bedrock of machine learning. You cannot develop a deep understanding and application of machine learning without it. Cut through the equations, Greek letters, and confusion, and discover the topics in probability that you need to know. Using clear explanations, standard Python libraries, and step-by-step tutorial lessons, you will discover the importance of probability to machine learning, Bayesian probability, entropy, density estimation, maximum likelihood, and much more.