

Deep Learning Adaptive Computation And Machine Learning Series

The fundamental mathematical tools needed to understand machine learning include linear algebra, analytic geometry, matrix decompositions, vector calculus, optimization, probability and statistics. These topics are traditionally taught in disparate courses, making it hard for data science or computer science students, or professionals, to efficiently learn the mathematics. This self-contained textbook bridges the gap between mathematical and machine learning texts, introducing the mathematical concepts with a minimum of prerequisites. It uses these concepts to derive four central machine learning methods: linear regression, principal component analysis, Gaussian mixture models and support vector machines. For students and others with a mathematical background, these derivations provide a starting point to machine learning texts. For those learning the mathematics for the first time, the methods help build intuition and practical experience with applying mathematical concepts. Every chapter includes worked examples and exercises to test understanding. Programming tutorials are offered on the book's web site. This book covers both classical and modern models in deep learning. The primary focus is on the theory and algorithms of deep learning. The theory and algorithms of neural networks are particularly important for understanding important concepts, so that one can understand the important design concepts of neural architectures in different applications. Why do neural networks work? When do they work better than off-the-shelf machine-learning models? When is depth useful? Why is training neural networks so hard? What are the pitfalls? The book is also rich in discussing different applications in order to give the practitioner a flavor of how neural architectures are designed for different types of problems. Applications associated with many different areas like recommender systems, machine translation, image captioning, image classification, reinforcement-learning based gaming, and text analytics are covered. The chapters of this book span three categories: The basics of neural networks: Many traditional machine learning models can be understood as special cases of neural networks. An emphasis is placed in the first two chapters on understanding the relationship between traditional machine learning and neural networks. Support vector machines, linear/logistic regression, singular value decomposition, matrix factorization, and recommender systems are shown to be special cases of neural networks. These methods are studied together with recent feature engineering methods like word2vec. Fundamentals of neural networks: A detailed discussion of training and regularization is provided in Chapters 3 and 4. Chapters 5 and 6 present radial-basis function (RBF) networks and restricted Boltzmann machines. Advanced topics in neural networks: Chapters 7 and 8 discuss recurrent neural networks and convolutional neural networks. Several advanced topics like deep reinforcement learning, neural Turing machines, Kohonen self-organizing maps, and generative adversarial networks are introduced in Chapters 9 and 10. The book is written for graduate students, researchers, and practitioners. Numerous exercises are available along with a solution manual to aid in classroom teaching. Where possible, an application-centric view is highlighted in order to provide an understanding of the practical uses of each class of techniques.

A comprehensive review of an area of machine learning that deals with the use of unlabeled data in classification problems: state-of-the-art algorithms, a taxonomy of the field, applications, benchmark experiments, and directions for future research. In the field of machine learning, semi-supervised learning (SSL) occupies the middle ground, between supervised learning (in which all training examples are labeled) and unsupervised learning (in which no label data are given). Interest in SSL has increased in recent years, particularly because of application domains in which unlabeled data are plentiful, such as images, text, and bioinformatics. This first comprehensive overview of SSL presents state-of-the-art algorithms, a taxonomy of the field, selected applications, benchmark experiments, and perspectives on ongoing and future research. Semi-Supervised Learning first presents the key assumptions and ideas underlying the field: smoothness, cluster or low-density separation, manifold structure, and transduction. The core of the book is the presentation of SSL methods, organized according to algorithmic strategies. After an examination of generative models, the book describes algorithms that implement the low-density separation assumption, graph-based methods, and algorithms that perform two-step learning. The book then discusses SSL applications and offers guidelines for SSL practitioners by analyzing the results of extensive benchmark experiments. Finally, the book looks at interesting directions for SSL research. The book closes with a discussion of the relationship between semi-supervised learning and transduction.

The book intends to cover various problematic aspects of emerging smart computing and

self-adapting technologies comprising of machine learning, artificial intelligence, deep learning, robotics, cloud computing, fog computing, data mining algorithms, including emerging intelligent and smart applications related to these research areas. Further coverage includes implementation of self-adaptation architecture for smart devices, self-adaptive models for smart cities and self-driven cars, decentralized self-adaptive computing at the edge networks, energy-aware AI-based systems, M2M networks, sensors, data analytics, algorithms and tools for engineering self-adaptive systems, and so forth. Acts as guide to Self-healing and Self-adaptation based fully automatic future technologies Discusses about Smart Computational abilities and self-adaptive systems Illustrates tools and techniques for data management and explains the need to apply, and data integration for improving efficiency of big data Exclusive chapter on the future of self-stabilizing and self-adaptive systems of systems Covers fields such as automation, robotics, medical sciences, biomedical and agricultural sciences, healthcare and so forth This book is aimed researchers and graduate students in machine learning, information technology, and artificial intelligence.

Machine Learning

Understanding Machine Learning

Dive Into Deep Learning

Encyclopedia of Machine Learning

Learning with Kernels

Introduction to Machine Learning

A concise and self-contained introduction to causal inference, increasingly important in data science and machine learning. The mathematization of causality is a relatively recent development, and has become increasingly important in data science and machine learning. This book offers a self-contained and concise introduction to causal models and how to learn them from data. After explaining the need for causal models and discussing some of the principles underlying causal inference, the book teaches readers how to use causal models: how to compute intervention distributions, how to infer causal models from observational and interventional data, and how causal ideas could be exploited for classical machine learning problems. All of these topics are discussed first in terms of two variables and then in the more general multivariate case. The bivariate case turns out to be a particularly hard problem for causal learning because there are no conditional independences as used by classical methods for solving multivariate cases. The authors consider analyzing statistical asymmetries between cause and effect to be highly instructive, and they report on their decade of intensive research into this problem. The book is accessible to readers with a background in machine learning or statistics, and can be used in graduate courses or as a reference for researchers. The text includes code snippets that can be copied and pasted, exercises, and an appendix with a summary of the most important technical concepts.

Use your Raspberry Pi to get smart about computing fundamentals In the 1980s, the tech revolution was kickstarted by a flood of relatively inexpensive, highly programmable computers like the Commodore. Now, a second revolution in computing is beginning with the Raspberry Pi. Learning Computer Architecture with the Raspberry Pi is the premier guide to understanding the components of the most exciting tech product available. Thanks to this book, every Raspberry Pi owner can understand how the computer works and how to access all of its hardware and software capabilities. Now, students, hackers, and casual users alike can discover how computers work with Learning Computer Architecture with the Raspberry Pi. This book explains what each and every hardware component does, how they relate to one another, and how they correspond to the components of other computing systems. You'll also learn how programming works and how the operating system relates to the Raspberry Pi's physical components. Co-authored by Eben Upton, one of the creators of the Raspberry Pi, this is a companion volume to the Raspberry Pi User Guide An affordable solution for learning about computer system design considerations and experimenting with low-level programming Understandable descriptions of the functions of memory storage, Ethernet, cameras, processors, and more Gain knowledge of computer design and operation in general by exploring the basic structure of the Raspberry Pi The Raspberry Pi was created to bring forth a new generation of computer scientists, developers, and architects who understand the inner workings of the computers that have become essential to our daily lives. Learning Computer Architecture with the Raspberry Pi is your gateway to the world of computer system design.

A rigorous and comprehensive textbook covering the major approaches to knowledge graphs, an active and interdisciplinary area within artificial intelligence. The field of knowledge graphs, which allows us to model, process, and derive insights from complex real-world data, has emerged as an active and interdisciplinary area of artificial intelligence over the last decade, drawing on such fields as natural language processing, data mining, and the semantic web. Current projects involve predicting cyberattacks, recommending products, and even gleaning insights from thousands of papers on COVID-19. This textbook offers rigorous and comprehensive coverage of the field. It focuses systematically on the major approaches, both those that have stood the test of time and the latest deep learning methods.

Although interest in machine learning has reached a high point, lofty expectations often scuttle projects before they get very far. How can machine learning—especially deep neural networks—make a real difference in your organization? This hands-on guide not only provides the most practical information available on the subject, but also helps you get started building efficient deep learning networks. Authors Adam Gibson and Josh Patterson provide theory on deep learning before introducing their open-source Deeplearning4j (DL4J) library for developing production-class workflows. Through real-world examples, you'll learn methods and strategies for training deep network architectures and running deep learning workflows on Spark and Hadoop with DL4J. Dive into machine learning concepts in general, as well as deep learning in particular Understand how deep networks evolved from neural network fundamentals Explore the major deep network architectures, including Convolutional and Recurrent Learn how to map specific deep networks to the right problem Walk through the fundamentals of tuning general neural networks and specific deep network architectures Use vectorization techniques for different data types with DataVec, DL4J's workflow tool Learn how to use DL4J natively on Spark and Hadoop

The Minimum Description Length Principle

The Machine Learning Approach

Learning Deep Learning

Foundations and Learning Algorithms

An Introduction

Theory and Practice of Neural Networks, Computer Vision, Nlp, and Transformers Using Tensorflow

A description of perturbation-based methods developed in machine learning to augment novel

optimization methods with strong statistical guarantees. In nearly all machine learning, decisions must be made given current knowledge. Surprisingly, making what is believed to be the best decision is not always the best strategy, even when learning in a supervised learning setting. An emerging body of work on learning under different rules applies perturbations to decision and learning procedures. These methods provide simple and highly efficient learning rules with improved theoretical guarantees. This book describes perturbation-based methods developed in machine learning to augment novel optimization methods with strong statistical guarantees, offering readers a state-of-the-art overview. Chapters address recent modeling ideas that have arisen within the perturbations framework, including Perturb & MAP, herding, and the use of neural networks to map generic noise to distribution over highly structured data. They describe new learning procedures for perturbation models, including an improved EM algorithm and a learning algorithm that aims to match moments of model samples to moments of data. They discuss understanding the relation of perturbation models to their traditional counterparts, with one chapter showing that the perturbations viewpoint can lead to new algorithms in the traditional setting. And they consider perturbation-based regularization in neural networks, offering a more complete understanding of dropout and studying perturbations in the context of deep neural networks.

The goal of machine learning is to program computers to use example data or past experience to solve a given problem. Many successful applications of machine learning exist already, including systems that analyze past sales data to predict customer behavior, optimize robot behavior so that a task can be completed using minimum resources, and extract knowledge from bioinformatics data. Introduction to Machine Learning is a comprehensive textbook on the subject, covering a broad array of topics not usually included in introductory machine learning texts. Subjects include supervised learning; Bayesian decision theory; parametric, semi-parametric, and nonparametric methods; multivariate analysis; hidden Markov models; reinforcement learning; kernel machines; graphical models; Bayesian estimation; and statistical testing. Machine learning is rapidly becoming a skill that computer science students must master before graduation. The third edition of Introduction to Machine Learning reflects this shift, with added support for beginners, including selected solutions for exercises and additional example data sets (with code available online). Other substantial changes include discussions of outlier detection; ranking algorithms for perceptrons and support vector machines; matrix decomposition and spectral methods; distance estimation; new kernel algorithms; deep learning in multilayered perceptrons; and the nonparametric approach to Bayesian methods. All learning algorithms are explained so that students can easily move from the equations in the book to a computer program. The book can be used by both advanced undergraduates and graduate students. It will also be of interest to professionals who are concerned with the application of machine learning methods. A modern treatment focusing on learning and inference, with minimal prerequisites, real-world examples and implementable algorithms.

NVIDIA's Full-Color Guide to Deep Learning: All Students Need to Get Started and Get Results

Learning Deep Learning is a complete guide to DL. Illuminating both the core concepts and the hands-on programming techniques needed to succeed, this book suits seasoned developers, data scientists, analysts, but also those with no prior machine learning or statistics experience.

After introducing the essential building blocks of deep neural networks, such as artificial neurons and fully connected, convolutional, and recurrent layers, Magnus Ekman shows how to use them to build advanced architectures, including the Transformer. He describes how these concepts are used to build modern networks for computer vision and natural language processing (NLP), including Mask R-CNN, GPT, and BERT. And he explains how a natural language translator and a system generating natural language descriptions of images. Throughout, Ekman provides concise, well-annotated code examples using TensorFlow with Keras. Corresponding PyTorch examples are provided online, and the book thereby covers the two dominating Python libraries for DL used in industry and academia. He concludes with an introduction to neural architecture search (NAS), exploring important ethical issues and providing resources for further learning. Explore and master core concepts: perceptrons, gradient-based learning, sigmoid neurons, and back propagation. See how DL frameworks make it easier to develop more complicated and useful neural networks. Discover how convolutional neural networks (CNNs) revolutionize image classification and analysis. Apply recurrent neural networks (RNNs) and long short-term memory (LSTM) to text and other variable-length sequences. Master NLP with sequence-to-sequence networks and the Transformer architecture. Build applications for natural language translation and image captioning.

Introduction to Deep Learning

Neural Networks and Deep Learning

Mathematics for Machine Learning

Fundamentals, Techniques, and Applications

Bioinformatics

Models, Learning, and Inference

A project-based guide to the basics of deep learning. This concise, project-driven guide to deep learning takes readers through a series of program-writing tasks that introduce them to the use of deep learning in such areas of artificial intelligence as computer vision, natural-language processing, and reinforcement learning. The author, a longtime artificial intelligence researcher specializing in natural-language processing, covers feed-forward neural nets, convolutional neural

nets, word embeddings, recurrent neural nets, sequence-to-sequence learning, deep reinforcement learning, unsupervised models, and other fundamental concepts and techniques. Students and practitioners learn the basics of deep learning by working through programs in Tensorflow, an open-source machine learning framework. "I find I learn computer science material best by sitting down and writing programs," the author writes, and the book reflects this approach. Each chapter includes a programming project, exercises, and references for further reading. An early chapter is devoted to Tensorflow and its interface with Python, the widely used programming language. Familiarity with linear algebra, multivariate calculus, and probability and statistics is required, as is a rudimentary knowledge of programming in Python. The book can be used in both undergraduate and graduate courses; practitioners will find it an essential reference.

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.

Deep Learning MIT Press

Get to grips with the basics of Keras to implement fast and efficient deep-learning models About This Book Implement various deep-learning algorithms in Keras and see how deep-learning can be used in games See how various deep-learning models and practical use-cases can be implemented using Keras A practical, hands-on guide with real-world examples to give you a strong foundation in Keras Who This Book Is For If you are a data scientist with experience in machine learning or an AI programmer with some exposure to neural networks, you will find this book a useful entry point to deep-learning with Keras. A knowledge of Python is required for this book. What You Will Learn Optimize step-by-step functions on a large neural network using the Backpropagation Algorithm Fine-tune a neural network to improve the quality of results Use deep learning for image and audio processing Use Recursive Neural Tensor Networks (RNTNs) to outperform standard word embedding in special cases Identify problems for which Recurrent Neural Network (RNN) solutions are suitable Explore the process required to implement Autoencoders Evolve a deep neural network using reinforcement learning In Detail This book starts by introducing you to supervised learning algorithms such as simple linear regression, the classical multilayer perceptron and more sophisticated deep convolutional networks. You will also explore image processing with recognition of hand written digit images, classification of images into different categories, and advanced objects recognition with related image annotations. An example of identification of salient points for face detection is also provided. Next you will be introduced to Recurrent Networks, which are optimized for processing sequence data such as text, audio or time series. Following that, you will learn about unsupervised learning algorithms such as Autoencoders and the very popular Generative Adversarial Networks (GAN). You will also explore non-traditional uses of neural networks as Style Transfer. Finally, you will look at Reinforcement Learning and its application to AI game playing, another popular direction of research and application of neural networks. Style and approach This book is an easy-to-follow guide full of examples and real-world applications to help you gain an in-depth understanding of Keras. This book will showcase more than twenty working Deep Neural Networks coded in Python using Keras.

Elements of Causal Inference

A Visual, Interactive Guide to Artificial Intelligence

Deep Learning with Keras

Artificial Intelligence

Deep Learning Illustrated

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

This comprehensive encyclopedia, in A-Z format, provides easy access to relevant information for those seeking entry into any aspect of Machine Learning. Most of the entries in this preeminent work include useful literature references.

A new edition of a graduate-level machine learning textbook that focuses on the analysis and theory of algorithms. This book is a general introduction to machine learning that can serve as a textbook for graduate students and a reference for researchers. It covers fundamental modern machine learning while providing the theoretical basis and conceptual tools needed for the discussion and justification of algorithms. It also describes various aspects of the application of these algorithms. The authors aim to present novel theoretical tools and concepts while giving concise practical advice on relatively advanced topics. Foundations of Machine Learning is unique in its focus on the analysis and theory of algorithms. The first part provides a theoretical foundation for what follows; subsequent chapters are mostly self-contained. Topics covered include the Probably Approximately Correct learning framework; generalization bounds based on Rademacher complexity and VC-dimension; Support Vector Machines (SVMs); kernel

boosting; on-line learning; multi-class classification; ranking; regression; algorithmic stability; dimensionality reduction; learning automata; languages; and reinforcement learning. Each chapter ends with a set of exercises. Appendixes provide additional material including concise review. This second edition offers three new chapters, on model selection, maximum entropy models, and conditional entropy models. New appendixes includes a major section on Fenchel duality, expanded coverage of concentration inequalities, and an entirely new entry on information theory. More than half of the exercises are new to this edition.

The leading experts in system change and learning, with their school-based partners around the world, have created this essential comp. runaway best-seller, *Deep Learning: Engage the World Change the World*. This hands-on guide provides a roadmap for building capacity in schools, districts, and systems to design deep learning, measure progress, and assess conditions needed to activate and sustain innovation. *Learning: Tools for Engagement* is rich with resources educators need to construct and drive meaningful deep learning experiences in our kind of mindset and know-how that is crucial to becoming a problem-solving change agent in our global society. Designed in full color, this guide is loaded with tools, tips, protocols, and real-world examples. It includes:

- A framework for deep learning that provides a pathway to the global competencies needed to flourish in a complex world — character, citizenship, collaboration, communication, creativity, and critical thinking.
- Learning progressions to help educators analyze student work and measure progress.
- Learning design rubrics, templates and examples of the four elements of learning design: learning partnerships, pedagogical practices, learning environments, and leveraging digital.
- Conditions for teacher self-assessment tools, and planning guides to help educators build, mobilize, and sustain deep learning in schools and districts.

improve, and expand your world of learning. Put the joy back into learning for students and adults alike. Dive into deep learning to create experiences that give purpose, unleash student potential, and transform not only learning, but life itself.

Introduction to Natural Language Processing

A Probabilistic Perspective

Reinforcement Learning, second edition

A Textbook

Deep Learning with Python, Second Edition

Deep Learning with PyTorch

A comprehensive introduction to Support Vector Machines and related kernel methods. In the 1990s, a new type of learning algorithm was developed, based on results from statistical learning theory: the Support Vector Machine (SVM). This gave rise to a new class of theoretically elegant learning machines that use a central concept of SVMs—kernels—for a number of learning tasks. Kernel machines provide a modular framework that can be adapted to different tasks and domains by the choice of the kernel function and the base algorithm. They are replacing neural networks in a variety of fields, including engineering, information retrieval, and bioinformatics. *Learning with Kernels* provides an introduction to SVMs and related kernel methods. Although the book begins with the basics, it also includes the latest research. It provides all of the concepts necessary to enable a reader equipped with some basic mathematical knowledge to enter the world of machine learning using theoretically well-founded yet easy-to-use kernel algorithms and to understand and apply the powerful algorithms that have been developed over the last few years.

“We finally have the definitive treatise on PyTorch! It covers the basics and abstractions in great detail. I hope this book becomes your extended reference document.” —Soumith Chintala, co-creator of PyTorch

Key Features Written by PyTorch’s creator and key contributors

- Develop deep learning models in a familiar Pythonic way
- Use PyTorch to build an image classifier for cancer detection
- Diagnose problems with your neural network and improve training with data augmentation

Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications.

About The Book Every other day we hear about new ways to put deep learning to good use: improved medical imaging, accurate credit card fraud detection, long range weather forecasting, and more. PyTorch puts these superpowers in your hands. Instantly familiar to anyone who knows Python data tools like NumPy and Scikit-learn, PyTorch simplifies deep learning without sacrificing advanced features. It’s great for building quick models, and it scales smoothly from laptop to enterprise. *Deep Learning with PyTorch* teaches you to create deep learning and neural network systems with PyTorch. This practical book gets you to work right away building a tumor image classifier from scratch. After covering the basics, you’ll learn best practices for the entire deep learning pipeline, tackling advanced projects as your PyTorch skills become more sophisticated. All code samples are easy to explore in downloadable Jupyter notebooks.

What You Will Learn

- Understanding deep learning data structures such as tensors and neural networks
- Best practices for the PyTorch Tensor API, loading data in Python, and visualizing results
- Implementing modules and loss functions
- Utilizing pretrained models from PyTorch Hub
- Methods for training networks with limited inputs
- Sifting through unreliable results to diagnose and fix problems in your neural network
- Improve your results with augmented data, better model architecture, and fine tuning

This Book Is Written For

For Python programmers with an interest in machine learning. No experience with PyTorch or other deep learning frameworks is required.

About The Authors

Eli Stevens has worked in Silicon Valley for the past 15 years as a software engineer, and the past 7 years as Chief Technical Officer of a startup making medical device software. Luca Antiga is co-founder and CEO of an AI engineering company located in Bergamo, Italy, and a regular contributor to PyTorch. Thomas Viehmann is a Machine Learning and PyTorch speciality trainer and consultant based in Munich, Germany and a PyTorch core developer.

Table of Contents

PART 1 - CORE PYTORCH

- 1 Introducing deep learning and the PyTorch Library
- 2 Pretrained networks
- 3 It starts with a tensor
- 4 Real-world data representation using tensors
- 5 The mechanics of learning
- 6 Using a neural network to fit the data
- 7 Telling birds from airplanes: Learning from images
- 8 Using convolutions to generalize

PART 2 - LEARNING FROM IMAGES IN THE REAL WORLD: EARLY DETECTION OF LUNG CANCER

- 9 Using PyTorch to fight cancer
- 10 Combining data sources into a unified dataset
- 11 Training a classification model to detect suspected tumors
- 12 Improving training with metrics and augmentation
- 13 Using segmentation to find suspected nodules
- 14 End-to-end nodule analysis, and where to go next

PART 3 - DEPLOYMENT

- 15 Deploying to production

This introduction to the MDL Principle provides a reference accessible to graduate students and researchers in statistics, pattern classification, machine learning, and data mining, to philosophers interested in the foundations of statistics, and to researchers in other applied sciences that involve model selection.

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

Semi-Supervised Learning

Principles and Techniques

with Practical Examples in MOA

Deep Learning

Knowledge Graphs

Probabilistic Graphical Models

A general framework for constructing and using probabilistic models of complex systems that would enable a computer to use available information for making decisions. Most tasks require a person or an automated system to reason—to reach conclusions based on available information. The framework of probabilistic graphical models, presented in this book, provides a general approach for this task. The approach is model-based, allowing interpretable models to be constructed and then manipulated by reasoning algorithms. These models can also be learned automatically from data, allowing the approach to be used in cases where manually constructing a model is difficult or even impossible. Because uncertainty is an inescapable aspect of most real-world applications, the book focuses on probabilistic models, which make the uncertainty explicit and provide models that are more faithful to reality. *Probabilistic Graphical Models* discusses a variety of models, spanning Bayesian networks, undirected Markov

networks, discrete and continuous models, and extensions to deal with dynamical systems and relational data. For each class of models, the text describes the three fundamental cornerstones: representation, inference, and learning, presenting both basic concepts and advanced techniques. Finally, the book considers the use of the proposed framework for causal reasoning and decision making under uncertainty. The main text in each chapter provides the detailed technical development of the key ideas. Most chapters also include boxes with additional material: skill boxes, which describe techniques; case study boxes, which discuss empirical cases related to the approach described in the text, including applications in computer vision, robotics, natural language understanding, and computational biology; and concept boxes, which present significant concepts drawn from the material in the chapter. Instructors (and readers) can group chapters in various combinations, from core topics to more technically advanced material, to suit their particular needs.

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.

A survey of computational methods for understanding, generating, and manipulating human language, which offers a synthesis of classical representations and algorithms with contemporary machine learning techniques. This textbook provides a technical perspective on natural language processing—methods for building computer software that understands, generates, and manipulates human language. It emphasizes contemporary data-driven approaches, focusing on techniques from supervised and unsupervised machine learning. The first section establishes a foundation in machine learning by building a set of tools that will be used throughout the book and applying them to word-based textual analysis. The second section introduces structured representations of language, including sequences, trees, and graphs. The third section explores different approaches to the representation and analysis of linguistic meaning, ranging from formal logic to neural word embeddings. The final section offers chapter-length treatments of three transformative applications of natural language processing: information extraction, machine translation, and text generation. End-of-chapter exercises include both paper-and-pencil analysis and software implementation. The text synthesizes and distills a broad and diverse research literature, linking contemporary machine learning techniques with the field's linguistic and computational foundations. It is suitable for use in advanced undergraduate and graduate-level courses and as a reference for software engineers and data scientists. Readers should have a background in computer programming and college-level mathematics. After mastering the material presented, students will have the technical skill to build and analyze novel natural language processing systems and to understand the latest research in the field.

A hands-on approach to tasks and techniques in data stream mining and real-time analytics, with examples in MOA, a popular freely available open-source software framework. Today many information sources—including sensor networks, financial markets, social networks, and healthcare monitoring—are so-called data streams, arriving sequentially and at high speed. Analysis must take place in real time, with partial data and without the capacity to store the entire data set. This book presents algorithms and techniques used in data stream mining and real-time analytics. Taking a hands-on approach, the book demonstrates the techniques using MOA (Massive Online Analysis), a popular, freely available open-source software framework, allowing readers to try out the techniques after reading the explanations. The book first offers a brief introduction to the topic, covering big data mining, basic methodologies for mining data streams, and a simple example of MOA. More detailed discussions follow, with chapters on sketching techniques, change, classification, ensemble methods, regression, clustering, and frequent pattern mining. Most of these chapters include exercises, an MOA-based lab session, or both. Finally, the book discusses the MOA software, covering the MOA graphical user interface, the command line, use of its API, and the development of new methods within MOA. The book will be an essential reference for readers who want to use data stream mining as a tool, researchers in innovation or data stream mining, and programmers who want to create new algorithms for MOA.

From Theory to Algorithms

Probabilistic Machine Learning

Cloud Computing for Machine Learning and Cognitive Applications

Perturbations, Optimization, and Statistics

Machine Learning for Data Streams

Computer Vision

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.

Unlock the groundbreaking advances of deep learning with this extensively revised edition of the bestselling original. Learn directly from the creator of Keras and master practical Python deep learning techniques that are easy to apply in the real world. In Deep Learning with Python, Second Edition you will learn: Deep learning from first principles Image classification & image segmentation Timeseries forecasting Text classification and machine translation Text generation, neural style transfer, and image generation Deep Learning with Python has taught thousands of readers how to put the full capabilities of deep learning into action. This extensively revised second edition introduces deep learning using Python and Keras, and is loaded with insights for both novice and experienced ML practitioners. You'll learn practical techniques that are easy to apply in the real world, and important theory for perfecting neural networks. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the technology Recent innovations in deep learning unlock exciting new software capabilities like automated language translation, image recognition, and more. Deep learning is becoming essential knowledge for every software developer, and modern tools like Keras and TensorFlow put it within your reach, even if you have no background in mathematics or data science. About the book Deep Learning with Python, Second Edition introduces the field of deep learning using Python and the powerful Keras library. In this new edition, Keras creator François Chollet offers insights for both novice and experienced machine learning practitioners. As you move through this book, you'll build your understanding through intuitive explanations, crisp illustrations, and clear examples. You'll pick up the skills to start developing deep-learning applications. What's inside Deep learning from first principles Image classification and image segmentation Time series forecasting Text classification and machine translation Text generation, neural style transfer, and image generation About the reader For readers with intermediate Python skills. No previous experience with Keras, TensorFlow, or machine learning is required. About the author François Chollet is a software engineer at Google and creator of the Keras deep-learning library. Table of Contents 1 What is deep learning? 2 The mathematical building blocks of neural networks 3 Introduction to Keras and TensorFlow 4 Getting started with neural networks: Classification and regression 5 Fundamentals of machine learning 6 The universal workflow of machine learning 7 Working with Keras: A deep dive 8 Introduction to deep learning for computer vision 9 Advanced deep learning for computer vision 10 Deep learning for timeseries 11 Deep learning for text 12 Generative deep learning 13 Best practices for the real world 14 Conclusions

An unprecedented wealth of data is being generated by genome sequencing projects and other experimental efforts to determine the structure and function of biological molecules. The demands and opportunities for interpreting these data are expanding more than ever.

Biotechnology, pharmacology, and medicine will be particularly affected by the new results and the increased understanding of life at the molecular level. Bioinformatics is the development and application of computer methods for analysis, interpretation, and prediction, as well as for the design of experiments. It has emerged as a strategic frontier between biology and computer science. Machine learning approaches (e.g., neural networks, hidden Markov models, and belief networks) are ideally suited for areas where there is a lot of data but little theory—and this is exactly the situation in molecular biology. As with its predecessor, statistical model fitting, the goal in machine learning is to extract useful information from a body of data by building good probabilistic models. The particular twist behind machine learning, however, is to automate the process as much as possible. In this book, Pierre Baldi and Soren Brunak present the key machine learning approaches and apply them to the computational problems encountered in the analysis of biological data. The book is aimed at two types of researchers and students. First are the biologists and biochemists who need to understand new data-driven algorithms, such as neural networks and hidden Markov models, in the context of biological sequences and their molecular structure and function. Second are those with a primary background in physics, mathematics, statistics, or computer science who need to know more about specific applications in molecular biology.

"Updated edition of popular textbook on Artificial Intelligence. This edition specific looks at ways of keeping artificial intelligence under control"--

A Modern Approach

Graphical Models for Machine Learning and Digital Communication

Foundations of Machine Learning, second edition

Introduction to Machine Learning, fourth edition

Gaussian Processes for Machine Learning

Smart Computing and Self-Adaptive Systems

A detailed and up-to-date introduction to machine learning, presented through the unifying lens of probabilistic modeling and Bayesian decision theory. This book offers a detailed and up-to-date introduction to machine learning (including deep learning) through the unifying lens of probabilistic modeling and Bayesian decision theory. The book covers mathematical background (including linear algebra and optimization), basic supervised learning (including linear and logistic regression and deep neural networks), as well as more advanced topics (including transfer learning and unsupervised learning). End-of-chapter exercises allow students to apply what they have learned, and an appendix covers notation. Probabilistic Machine Learning grew out of the author's 2012 book, Machine Learning: A Probabilistic Perspective. More than just a simple update, this is a completely new book that reflects the dramatic developments in the field since 2012, most notably deep learning. In addition, the new book is accompanied by online Python code, using libraries such as scikit-learn, JAX, PyTorch, and Tensorflow, which can be used to reproduce nearly all the figures; this code can be run inside a web browser using cloud-based notebooks, and provides a practical complement to the theoretical topics discussed in the book. This introductory text will be followed by a sequel that covers more advanced topics, taking the same probabilistic approach.

An accessible introduction to the artificial intelligence technology that enables computer vision, speech recognition, machine translation, and driverless cars. Deep learning is an artificial intelligence technology that enables computer vision, speech recognition in mobile phones, machine translation, AI games, driverless cars, and other applications. When we use consumer products from Google, Microsoft, Facebook, Apple, or Baidu, we are often interacting with a deep learning system. In this volume in the MIT Press Essential Knowledge series, computer scientist John Kelleher offers an accessible and concise but comprehensive introduction to the fundamental technology at the heart of the artificial intelligence revolution. Kelleher explains that deep learning enables data-driven decisions by identifying and extracting patterns from large datasets; its ability to learn from complex data makes deep learning ideally suited to take advantage of the rapid growth in big data and computational power. Kelleher also explains some of the basic concepts in deep learning, presents a history of advances in the field, and discusses the current state of the art. He describes the most important deep learning architectures, including autoencoders, recurrent neural networks, and long short-term networks, as well as such recent developments as Generative Adversarial Networks and capsule networks. He also provides a comprehensive (and comprehensible) introduction to the two fundamental algorithms

in deep learning: gradient descent and backpropagation. Finally, Kelleher considers the future of deep learning—major trends, possible developments, and significant challenges.

The first textbook to teach students how to build data analytic solutions on large data sets using cloud-based technologies. This is the first textbook to teach students how to build data analytic solutions on large data sets (specifically in Internet of Things applications) using cloud-based technologies for data storage, transmission and mashup, and AI techniques to analyze this data. This textbook is designed to train college students to master modern cloud computing systems in operating principles, architecture design, machine learning algorithms, programming models and software tools for big data mining, analytics, and cognitive applications. The book will be suitable for use in one-semester computer science or electrical engineering courses on cloud computing, machine learning, cloud programming, cognitive computing, or big data science. The book will also be very useful as a reference for professionals who want to work in cloud computing and data science. Cloud and Cognitive Computing begins with two introductory chapters on fundamentals of cloud computing, data science, and adaptive computing that lay the foundation for the rest of the book. Subsequent chapters cover topics including cloud architecture, mashup services, virtual machines, Docker containers, mobile clouds, IoT and AI, inter-cloud mashups, and cloud performance and benchmarks, with a focus on Google's Brain Project, DeepMind, and X-Lab programs, IBKai HwangM SyNapse, Bluemix programs, cognitive initiatives, and neurocomputers. The book then covers machine learning algorithms and cloud programming software tools and application development, applying the tools in machine learning, social media, deep learning, and cognitive applications. All cloud systems are illustrated with big data and cognitive application examples.

A substantially revised fourth edition of a comprehensive textbook, including new coverage of recent advances in deep learning and neural networks. The goal of machine learning is to program computers to use example data or past experience to solve a given problem. Machine learning underlies such exciting new technologies as self-driving cars, speech recognition, and translation applications. This substantially revised fourth edition of a comprehensive, widely used machine learning textbook offers new coverage of recent advances in the field in both theory and practice, including developments in deep learning and neural networks. The book covers a broad array of topics not usually included in introductory machine learning texts, including supervised learning, Bayesian decision theory, parametric methods, semiparametric methods, nonparametric methods, multivariate analysis, hidden Markov models, reinforcement learning, kernel machines, graphical models, Bayesian estimation, and statistical testing. The fourth edition offers a new chapter on deep learning that discusses training, regularizing, and structuring deep neural networks such as convolutional and generative adversarial networks; new material in the chapter on reinforcement learning that covers the use of deep networks, the policy gradient methods, and deep reinforcement learning; new material in the chapter on multilayer perceptrons on autoencoders and the word2vec network; and discussion of a popular method of dimensionality reduction, t-SNE. New appendixes offer background material on linear algebra and optimization. End-of-chapter exercises help readers to apply concepts learned. Introduction to Machine Learning can be used in courses for advanced undergraduate and graduate students and as a reference for professionals.

A Practitioner's Approach

Tools for Engagement

Learning Computer Architecture with Raspberry Pi

Support Vector Machines, Regularization, Optimization, and Beyond

Content Description. #Includes bibliographical references and index.

"The authors' clear visual style provides a comprehensive look at what's currently possible with artificial neural networks as well as a glimpse of the magic that's to come." -Tim Urban, author of Wait But Why Fully Practical, Insightful Guide to Modern Deep Learning Deep learning is transforming software, facilitating powerful new artificial intelligence capabilities, and driving unprecedented algorithm performance. Deep Learning Illustrated is uniquely intuitive and offers a complete introduction to the discipline's techniques. Packed with full-color figures and easy-to-follow code, it sweeps away the complexity of building deep learning models, making the subject approachable and fun to learn. World-class instructor and practitioner Jon Krohn-with visionary content from Grant Beyleveld and beautiful illustrations by Aglaé Bassens-presents straightforward analogies to explain what deep learning is, why it has become so popular, and how it relates to other machine learning approaches. Krohn has created a practical reference and tutorial for developers, data scientists, researchers, analysts, and students who want to start applying it. He illuminates theory with hands-on Python code in accompanying Jupyter notebooks. To help you progress quickly, he focuses on the versatile deep learning library Keras to nimbly construct efficient TensorFlow models; PyTorch, the leading alternative library, is also covered. You'll gain a pragmatic understanding of all major deep learning approaches and their uses in applications ranging from machine vision and natural language processing to image generation and game-playing algorithms. Discover what makes deep learning systems unique, and the implications for practitioners Explore new tools that make deep learning models easier to build, use, and improve Master essential theory: artificial neurons, training, optimization, convolutional nets, recurrent nets, generative adversarial networks (GANs), deep reinforcement learning, and more Walk through building interactive deep learning applications, and move forward with your own artificial intelligence projects Register your book for convenient access to downloads, updates, and/or corrections as they become available. See inside book for details.