

Deep Learning A Practitioners Approach

Practical Machine Learning for Data Analysis Using Python is a problem solver's guide for creating real-world intelligent systems. It provides a comprehensive approach with concepts, practices, hands-on examples, and sample code. The book teaches readers the vital skills required to understand and solve different problems with machine learning. It teaches machine learning techniques necessary to become a successful practitioner, through the presentation of real-world case studies in Python machine learning ecosystems. The book also focuses on building a foundation of machine learning knowledge to solve different real-world case studies across various fields, including biomedical signal analysis, healthcare, security, economics, and finance. Moreover, it covers a wide range of machine learning models, including regression, classification, and forecasting. The goal of the book is to help a broad range of readers, including IT professionals, analysts, developers, data scientists, engineers, and graduate students, to solve their own real-world problems. Offers a comprehensive overview of the application of machine learning tools in data analysis across a wide range of subject areas Teaches readers how to apply machine learning techniques to biomedical signals, financial data, and healthcare data Explores important classification and regression algorithms as well as other machine learning techniques Explains how to use Python to handle data extraction, manipulation, and exploration techniques, as well as how to visualize data spread across multiple dimensions and extract useful features

Engage the World Change the World Deep Learning has claimed the attention of educators and policymakers around the world. This book not only defines what deep learning is, but takes up the question of how to mobilize complex, whole-system change and transform learning for all students. Deep Learning is a global partnership that works to: transform the role of teachers to that of activators who design experiences that build global competencies using real-life problem solving; and supports schools, districts, and systems to shift practice and how to measure learning in authentic ways. This comprehensive strategy incorporates practical tools and processes to engage students, educators, and families in new partnerships and drive deep learning.

Deep Learning and Parallel Computing Environment for Bioengineering Systems delivers a significant forum for the technical advancement of deep learning in parallel computing environment across bio-engineering diversified domains and its applications. Pursuing an interdisciplinary approach, it focuses on methods used to identify and acquire valid, potentially useful knowledge sources. Managing the gathered knowledge and applying it to multiple domains including health care, social networks, mining, recommendation systems, image processing, pattern recognition and predictions using deep learning paradigms is the major strength of this book. This book integrates the core ideas of deep learning and its applications in bio engineering application domains, to be accessible to all scholars and academicians. The proposed techniques and concepts in this book can be extended in future to accommodate changing business organizations' needs as well as practitioners' innovative ideas. Presents novel, in-depth research contributions from a methodological/application perspective in understanding the fusion of deep machine learning paradigms and their capabilities in solving a diverse range of problems Illustrates the state-of-the-art and recent developments in the new theories and applications of deep learning approaches applied to parallel computing environment in bioengineering systems Provides concepts and technologies that are successfully used in the implementation of today's intelligent data-centric critical systems and multi-media Cloud/Big data

The present book is primarily intended for undergraduate and postgraduate students of computer science and engineering, information technology, and electrical and electronics engineering. It bridges the gaps in knowledge of the seemingly difficult areas of machine learning and nature inspired computing. The text is written in a highly interactive manner, which satisfies the learning curiosity of any reader. Content of the text has been diligently organized to offer seamless learning experience. The text begins with introduction to machine learning, which is followed by explanation of different aspects of machine learning. Various supervised, unsupervised, reinforced and nature inspired learning techniques are included in the text book with numerous examples and case studies. Different aspects of new machine learning and nature inspired learning algorithms are explained in-depth. The well-explained algorithms and pseudo codes for each topic make this book useful for students. The book also throws light on areas like prediction and classification systems. Key Features • Day to day examples and pictorial representations for deeper understanding of the subject • Helps readers easily create programs/applications • Research oriented approach • More case studies and worked-out examples for each machine learning algorithm than any other book

Machine Learning

Engage the World Change the World

Deep Learning and Parallel Computing Environment for Bioengineering Systems

Explainable AI with Python

A Practitioner's Approach

Techniques, Approaches, and Applications

Deep learning is often viewed as the exclusive domain of math PhDs and big tech companies. But as this hands-on guide demonstrates, programmers comfortable with Python can achieve impressive results in deep learning with little math background, small amounts of data, and minimal code. How? With fastai, the first library to provide a most frequently used deep learning applications. Authors Jeremy Howard and Sylvain Gugger, the creators of fastai, show you how to train a model on a wide range of tasks using fastai and PyTorch. You'll also dive progressively further into deep learning theory to gain a complete understanding of the algorithms behind the scenes. Train models on natural language processing, tabular data, and collaborative filtering Learn the latest deep learning techniques that matter most in practice Improve accuracy, speed, and reliability by understanding how deep learning models work Discover how to turn your models into web applications Implement deep learning algorithms from scratch Com your work Gain insight from the foreword by PyTorch cofounder, Soumith Chintala

This book provides a full presentation of the current concepts and available techniques to make "machine learning" systems more explainable. The approaches presented can be applied to almost all the current "machine learning" models: linear and logistic regression, deep learning neural networks, natural language processing and image recognition. Progress in Machine Learning is increasing the use of artificial agents to perform critical tasks previously handled by humans (healthcare, legal and finance, among others). While the principles that guide the design of these agents are understood, most of the current deep-learning models are "opaque" to human understanding. Exploring current gap in literature on this emerging topic by taking both a theoretical and a practical perspective, making the reader quickly capable of working with tools and code for Explainable AI. Beginning with examples of what Explainable AI (XAI) is and why it is needed in the field, the book details different approaches to XAI depending on whether you work on interpretable models with specific examples leveraging Python are then presented, showing how intrinsic interpretable models can be interpreted and how to produce "human understandable" explanations. Model-agnostic methods for XAI are shown to produce explanations without relying on ML models internals that are "opaque" to Computer Vision, the authors then look at explainable models for Deep Learning and prospective methods for the future. Taking a practical perspective, the authors demonstrate how to effectively use ML and XAI in science. The final chapter explains Adversarial Machine Learning and how to do XAI with adversarial examples.

Learn how to solve challenging machine learning problems with TensorFlow, Google's revolutionary new software library for deep learning. If you have some background in basic linear algebra and calculus, this practical book introduces machine-learning fundamentals by showing you how to design systems capable of detecting objects in images, analyzing video, and predicting the properties of potential medicines. TensorFlow for Deep Learning teaches concepts through practical examples and helps you build knowledge of deep learning foundations from the ground up. It's ideal for practicing developers with experience designing software systems, and useful for scientists and other scripting but not necessarily with designing learning algorithms. Learn TensorFlow fundamentals, including how to perform basic computation Build simple learning systems to understand their mathematical foundations Dive into fully connected deep networks used in thousands of applications Turn prototypes into high-quality models with Process images with convolutional neural networks Handle natural language datasets with recurrent neural networks Use reinforcement learning to solve games such as tic-tac-toe Train deep networks with hardware including GPUs and tensor processing units

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

Building with Python from First Principles

Designing Next-Generation Machine Intelligence Algorithms

Introduction to Deep Learning

Deep Learning with Python

The Deep Learning Revolution

DL4J and Beyond

Deep Learning for Robot Perception and Cognition introduces a broad range of topics and methods in deep learning for robot perception and cognition together with end-to-end methodologies. The book provides the conceptual and mathematical background needed for approaching a large number of robot perception and cognition tasks from an end-to-end learning point-of-view. The book is suitable for students, university and industry researchers and practitioners in Robotic Vision, Intelligent Control, Mechatronics, Deep Learning, Robotic Perception and Cognition tasks. Presents deep learning principles and methodologies Explains the principles of applying end-to-end learning in robotics applications Present how to design and train deep learning models Shows how to apply deep learning in robot vision tasks such as object recognition, image classification, video analysis, and more Uses robot simulation environments for training deep learning models Applies deep learning methods for different tasks ranging from planning and navigation to biological analysis

How deep learning—from Google Translate to driverless cars to personal cognitive assistants—is changing our lives and transforming every sector of the economy. The deep learning revolution has brought us driverless cars, the greatly improved Google Translate, fluent conversations with Siri and Alexa, and enormous profits from automated trading on the New York Stock Exchange. Deep learning networks can play poker better than professional poker players and defeat a world champion at Go. In this book, Terry Sejnowski explains how deep learning went from being an arcane academic field to a disruptive technology in the information economy. Sejnowski played an important role in the founding of deep learning, as one of a small group of researchers in the 1980s who challenged the prevailing logic-and-symbol based version of AI. The new version of AI Sejnowski and others developed, which became deep learning, is fueled instead by data. Deep networks learn from data in the same way that babies explore the world, starting with fresh eyes and gradually acquiring the skills needed to navigate novel environments. Learning algorithms extract information from raw data; information can be used to create knowledge; knowledge underlies understanding; understanding leads to wisdom. Someday a driverless car will know the road better than you do and drive with more skill; a deep learning network will diagnose your illness; a personal cognitive assistant will augment your puny human brain. It took nature many millions of years to evolve human intelligence; AI is on a trajectory measured in decades. Sejnowski prepares us for a deep learning future.

Implement natural language processing applications with Python using a problem-solution approach. This book has numerous coded exercises that will help you to quickly deploy natural language processing techniques, such as text classification, parts of speech identification, topic modeling, text summarization, text generation, and sentiment analysis. Natural Language Processing Recipes starts by offering solutions for cleaning and preprocessing text data and ways to analyze it with advanced algorithms. You'll see practical applications of the semantic as well as syntactic analysis of text, as well as complex natural language processing approaches that involve text normalization, advanced preprocessing, POS tagging, and sentiment analysis. You will also learn various applications of machine learning and deep learning in natural language processing. By using the recipes in this book, you will have a toolbox of solutions to apply to your own projects in the real world, making your development time quicker and more efficient. What You Will Learn Apply NLP techniques using Python libraries such as NLTK, TextBlob, spaCy, Stanford CoreNLP, and many more Implement the concepts of information retrieval, text summarization, sentiment analysis, and other advanced natural language processing techniques. Identify machine learning and deep learning techniques for natural language processing and natural language generation problems Who This Book Is For Data scientists who want to refresh and learn various concepts of natural language processing through coded exercises.

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 techniquesneeded to succeed, this book suits seasoned developers, data scientists, analysts, but also those with no prior machine learning or statisticsexperience. 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, includingthe Transformer. He describes how these concepts are used to build modernnetworks for computer vision and natural language processing (NLP), includingMask R-CNN, GPT, and BERT. And he explains how a natural language translatorat a system generating natural language descriptions of images. Throughout, Ekman provides concise, well-annotated code examples usingTensorFlow with Keras. Corresponding PyTorch examples are provided online, andthe book thereby covers the two dominating Python libraries for DL, due to industry and academia. He concludes with an introduction to neural architecturesearch (NAS), exploring important ethical issues and providing resources forfurther learning. Exploreand master core concepts: perceptrons, gradient-based learning, sigmoidneurons, and back propagation See how DL frameworks make it easier to developmore complicated and useful neural networks Discover how convolutional neuralnetworks (CNNs) revolutionize image classification and analysis Apply recurrentneural networks (RNNs) and long short-term memory (LSTM) to text and othervariable-length sequences Master NLP with sequence-to-sequence networks and theTransformer architecture Build applications for natural language translation andimage captioning

Deep Learning with PyTorch

A Probabilistic Perspective

Deep Learning for Biomedical Data Analysis

Interpretable Machine Learning

A Practical Implementation Guide to Predictive Data Analytics Using Python

Unlocking Text Data with Machine Learning and Deep Learning using Python

Deep LearningDL4J and BeyondO'Reilly Media

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—PMK (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.

Learn how to apply test-driven development (TDD) to machine-learning algorithms—and catch mistakes that could sink your analysis. In this practical guide, author Matthew Kirk takes you through the principles of TDD and machine learning, and shows you how to apply TDD to several machine-learning algorithms, including Naive Bayesian Classifiers and Neural Networks. Machine-learning algorithms often have tests baked in, but they can't account for human errors in coding. Rather than blindly rely on machine-learning results as many researchers have, you can mitigate the risk of errors with TDD and write clean, stable machine-learning code. If you're familiar with Ruby 2.1, you're ready to start. Apply TDD to write and run tests before you start coding Learn the best uses and tradeoffs of eight machine learning algorithms Use real-world examples to test each algorithm through engaging, hands-on exercises Understand the similarities between TDD and the scientific method for validating solutions Be aware of the risks of machine learning, such as underfitting and overfitting data Explore techniques for improving your machine-learning models or data extraction

Summary Deep Learning with Python introduces the field of deep learning using the Python language and the powerful Keras library. Written by Keras creator and Google AI researcher François Chollet, this book builds your understanding through intuitive explanations and practical examples. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the technology Machine learning has made remarkable progress in recent years. We went from near-unsusable speech and image recognition, to near-human, researchers. We went from a world champion, to defeating a world champion. Behind this progress is deep learning—a combination of engineering advances, best practices, and theory that enables a wealth of previously impossible smart applications. About the Book Deep Learning with Python introduces the field of deep learning using the Python language and the powerful Keras library. Written by Keras creator and Google AI researcher François Chollet, this book builds your understanding through intuitive explanations and practical examples. You'll explore challenging concepts and practice with applications in computer vision, natural-language processing, and generative models. By the time you finish, you'll have the knowledge and hands-on skills to apply deep learning in your own projects. What's Inside Deep learning from first principles Setting up your own deep-learning environment Image-classification models Deep learning for text and sequences Neural style transfer, text generation, and image generation About the Reader Readers need intermediate Python skills. No previous experience with Keras, TensorFlow, or machine learning is required. About the Author François Chollet works on deep learning at Google in Mountain View, CA. He is the creator of the Keras deep-learning library, as well as a contributor to the TensorFlow machine-learning framework. He also does deep-learning research, with a focus on computer vision and the application of machine learning to formal reasoning. His papers have been published at major conferences in the field, including the Conference on Computer Vision and Pattern Recognition (CVPR), the Conference and Workshop on Neural Information Processing Systems (NIPS), the International Conference on Artificial Intelligence and Statistics (ICML), and others. Table of Contents PART I - FUNDAMENTALS OF DEEP LEARNING What is deep learning? Before we begin: the mathematical building blocks of neural networks Getting started with neural networks Fundamentals of machine learning PART 2 - DEEP LEARNING IN PRACTICE Deep learning for computer vision Deep learning for text and sequences Advanced deep-learning best practices Generative deep learning Conclusions appendix A - Installing Keras and its dependencies on Ubuntu appendix B - Running Jupyter notebooks on an EC2 GPU instance

Deep Learning

A Hands-on Introduction

Tools for Engagement

Scaling Up Machine Learning

Deep Learning with Python, Second Edition

Learning Deep Learning

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

This book is dedicated to addressing the major challenges in fighting COVID-19 using artificial intelligence (AI) and machine learning (ML) – from cost and complexity to availability and accuracy. The aim of this book is to focus on both the design and implementation of AI-based approaches in proposed COVID-19 solutions that are enabled and supported by sensor networks, cloud computing, and 5G and beyond. This book presents research that contributes to the application of ML techniques to the problem of computer communication-assisted diagnosis of COVID-19 and similar diseases. The authors present the latest theoretical developments, real-world applications, and future perspectives on this topic. This book brings together a broad multidisciplinary community, aiming to integrate ideas, theories, models, and techniques from across different disciplines on intelligent solutions/systems, and to inform how cognitive systems in Next Generation Networks (NGN) should be designed, developed, and evaluated while exchanging and processing critical health information. Targeted readers are from varying disciplines who are interested in implementing the smart planet/environments vision via wireless/wired enabling technologies.

Roughly inspired by the human brain, deep neural networks trained with large amounts of data can solve complex tasks with unprecedented accuracy. This practical book provides an end-to-end guide to TensorFlow, the leading open source software library that helps you build and train neural networks for computer vision, natural language processing (NLP), speech recognition, and general predictive analytics. Authors Tom Hope, Yehezkel Resheff, and Itay Lieder provide a hands-on approach to TensorFlow fundamentals for a broad technical audience—from data scientists and engineers to students and researchers. You'll begin by working through some basic examples in TensorFlow before diving deeper into topics such as neural network architectures, TensorBoard visualization, TensorFlow abstraction libraries, and multithreaded input pipelines. Once you finish this book, you'll know how to build and deploy production-ready deep learning systems in TensorFlow. Get up and running with TensorFlow, rapidly and painlessly Learn how to use TensorFlow to build deep learning models from the ground up Train popular deep learning models for computer vision and NLP Use extensive abstraction libraries to make development easier and faster Learn how to scale TensorFlow, and use clusters to distribute model training Deploy TensorFlow in a production setting

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, 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

Deep Learning for Robot Perception and Cognition

From Linear Regression to Reinforcement Learning

Practical Recipes to Get Started Quickly

Machine Learning for Risk Calculations

Thoughtful Machine Learning

Practical Machine Learning for Data Analysis Using Python

The leading experts in system change and learning, with their school-based partners around the world, have created this essential companion to their runaway best-seller, Deep Learning: Engage the World Change the World. This hands-on guide provides a roadmap for building capacity in teachers, schools, districts, and systems to design deep learning, measure progress, and assess conditions needed to activate and sustain innovation. Dive into Deep Learning: Tools for Engagement is rich with resources educators need to construct and drive meaningful deep learning experiences in order to develop the 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 easy-to-use guide is loaded with tools, tips, protocols, and real-world examples. It includes: • A framework for deep learning that provides a pathway to develop the six 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 for incorporating the four elements of learning design: learning partnerships, pedagogical practices, learning environments, and leveraging digital. • Conditions rubrics, teacher self-assessment tools, and planning guides to help educators build, mobilize, and sustain deep learning in schools and districts. Learn about, improve, and expand your world of learning. Put the joy back into learning for students and adults alike. Dive into deep learning to create learning experiences that give purpose, unleash student potential, and transform not only learning, but life itself.

Deep learning doesn't have to be intimidating. Until recently, this machine-learning method required years of study, but with frameworks such as Keras and Tensorflow, software engineers without a background in machine learning can quickly enter the field. With the recipes in this cookbook, you'll learn how to solve deep-learning problems for classifying and generating text, images, and music. Each chapter consists of several recipes needed to complete a single project, such as training a music recommending system. Author Douwe Osinga also provides a chapter with half a dozen techniques to help you if you're stuck. Examples are written in Python with code available on GitHub as a set of Python notebooks. You'll learn how to: Create applications that will serve real users Use word embeddings to calculate text similarity Build a movie recommender system based on Wikipedia links Learn how AIs see the world by visualizing their internal state Build a model to suggest emojis for pieces of text Reuse pretrained networks to build an inverse image search service Compare how GANs, autoencoders and LSTMs generate icons Detect music styles and index song collections

This is the first rigorous, self-contained treatment of the theory of deep learning. Starting with the foundations of the theory and building it up, this is essential reading for any scientists, instructors, and students interested in artificial intelligence and deep learning. It provides guidance on how to think about scientific questions, and leads readers through the history of the field and its fundamental connections to neuroscience. The author discusses many applications to beautiful problems in the natural sciences, in physics, chemistry, and biomedicine. Examples include the search for exotic particles and dark matter in experimental physics, the prediction of molecular properties and reaction outcomes in chemistry, and the prediction of protein structures and the diagnostic analysis of biomedical images in the natural sciences. The text is accompanied by a full set of exercises at different difficulty levels and encourages out-of-the-box thinking.

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 sense of the practical use of different types of problems. Applications include systems like recommender systems, machine translation, image captioning, image classification, reinforcement learning based gaming, and text analysis 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.

Artificial Intelligence and Machine Learning for COVID-19

Deep Learning Illustrated

Parallel and Distributed Approaches

A Constraint-Based Approach

A Textbook

Deep Learning from Scratch

Looking for one central source where you can learn key findings on machine learning? Deep Learning: The Definitive Guide provides developers and data scientists with the most practical information available on the subject, including deep learning theory, best practices, and use cases. Authors Adam Gibson and Josh Patterson present the latest relevant papers and techniques in a nonacademic manner, and implement the core mathematics in their DL4J library. If you work in the embedded, desktop, and big data/Hadoop spaces and really want to understand deep learning, this is your book. "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 artificial 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 Beyelved and beautiful illustrations by Aglae 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 illustrate the concepts, but also includes the best practices of implementing them, and the software which can be used. Review the fundamentals of deep learning and Chebyshev tensors Discover pioneering algorithmic techniques that can create new opportunities in complex risk calculation Learn how to apply the solutions to a wide range of real-life risk calculations. Download sample code used in the book, so you can follow along and experiment with your own calculations Realize improved risk management whilst overcoming the burden of limited computational power Quants, IT professionals, and financial risk managers will benefit from this practitioner-oriented approach to state-of-the-art risk calculation.

With the reinvention of neural networks in the 2000s, deep learning has become an extremely active area of research, one that's paving the way for modern machine learning. In this practical book, author Nikhil Buduma provides examples and clear explanations to guide you through major concepts of this complicated field. Companies such as Google, Microsoft, and Facebook are actively growing in-house deep-learning teams. For the rest of us, however, deep learning is still a pretty complex and difficult subject to grasp. If you're familiar with Python, and have a background in calculus, along with a basic understanding of machine learning, this book will get you started. Examine the foundations of machine learning and neural networks Learn how to train feed-forward neural networks Use TensorFlow to implement your first neural network Manage problems that arise as you begin to make networks deeper Build neural networks that analyze complex images Perform effective dimensionality reduction using autoencoders Dive deep into sequence analysis to examine language Learn the fundamentals of reinforcement learning

A PRACTITIONER'S APPROACH

Natural Language Processing Recipes

Neural Networks and Deep Learning

A Test-Driven Approach

A Practitioner's View

Empirical Approach to Machine Learning

With the resurgence of neural networks in the 2010s, deep learning has become essential for machine learning practitioners and even many software engineers. This book provides a comprehensive introduction for data scientists and software engineers with machine learning experience. You'll start with deep learning basics and move quickly to the state-of-the-art advanced architectures, implementing everything from scratch along the way. Author Seth Weidman shows you how neural networks work using a first principles approach. You'll learn how to apply multilayer neural networks, convolutional neural networks, and recurrent neural networks from the ground up. With a thorough understanding of how neural networks work mathematically, computationally, and conceptually, you'll be set up for success on all future deep learning projects. This book provides: Extremely clear and thorough mental models—accompanied by working code examples and mathematical explanations—for understanding neural networks Methods for implementing multilayer neural networks from scratch, using an easy-to-understand object-oriented framework Working implementations and clear-cut explanations of convolutional and recurrent neural networks Implementation of these neural network concepts using the popular PyTorch framework Whether you're a software engineer aspiring to enter the world of deep learning, a veteran data scientist, or a hobbyist with a simple dream of making the next viral AI app, you might have wondered where to begin. This step-by-step guide teaches you how to build practical deep learning applications for the cloud, mobile, browsers, and edge devices using a hands-on approach. Relying on years of industry experience transforming deep learning research into award-winning applications, Anirudh Koul, Siddha Ganju, and Meher Kasam guide you through the process of converting an idea into something that people in the real world can use. Train, tune, and deploy computer vision models with Keras, TensorFlow, Core ML, and TensorFlow Lite Develop AI for a range of devices including Raspberry Pi, Jetson Nano, and Google Coral Explore fun projects, from Silicon Valley's Not Hotdog app to 40+ industry case studies Simulate an autonomous car in a video game environment and build a miniature version with reinforcement learning Use transfer learning to train models in minutes Discover 50+ practical tips for maximizing model accuracy and speed, debugging, and scaling to millions of users

Deep learning neural networks have become easy to define and fit, but are still hard to configure. Discover exactly how to improve the performance of deep learning neural network models on your predictive modeling projects. With clear explanations, standard Python libraries, and step-by-step tutorial lessons, you'll discover how to better train your models, reduce overfitting, and make more accurate predictions. The book explores various applications of deep learning to the diagnosis of cancer, while also outlining the future face of deep learning-assisted cancer diagnostics. As is commonly known, artificial intelligence has paved the way for countless new solutions in the field of medicine. In this context, deep learning is a recent and remarkable sub-field, which can effectively handle huge amounts of data and deliver more accurate results. As a vital research area, medical diagnosis is among those which deep learning-oriented solutions are often employed. Accordingly, the objective of this book is to highlight recent advanced applications of deep learning for diagnosing different types of cancer. The target audience includes scientists, experts, MSc and PhD students, postdocs, and anyone interested in the subjects discussed. The book can be used as a reference work to support courses on artificial intelligence, medical and biomedical education.

A Visual, Interactive Guide to Artificial Intelligence

Deep Learning in Science

TensorFlow for Deep Learning

Deep Learning for Coders with fastai and PyTorch

Deep Learning Cookbook

Fundamentals of Deep Learning

This integrated collection covers a range of parallelization platforms, concurrent programming frameworks and machine learning settings, with case studies.

This book is the first comprehensive Deep Learning DL for medical data analysis. It surveys the most recent techniques and approaches in this field, with both a broad coverage and enough depth to be of practical use to working professionals. This book offers enough fundamental and technical information on these techniques, approaches and the related problems without overcredoring the reader's head. It presents the results of the latest research in the field of DL for biomedical data analysis. The techniques and approaches presented in this book deal with the most important and/or the newest topics encountered in this field. They combine fundamental theory of Artificial Intelligence (AI), Machine Learning (ML) and DL with practical applications in Biology and Medicine. Certainly, the list of topics covered in this book is not exhaustive but these topics will shed light on the implications of the presented techniques and approaches on other topics in biomedical data analysis. The book finds a balance between theoretical and practical coverage of a wide range of issues in the field of biomedical data analysis, thanks to DL. The few published books on DL for biomedical data analysis either focus on specific topics or lack technical depth. The chapters presented in this book were selected for quality and relevance. The book also presents experiments that provide qualitative and quantitative overviews in the field of biomedical data analysis. The reader will require some familiarity with AI, ML and DL and will learn about techniques and approaches that deal with the most important and/or the newest topics encountered in the field of DL for biomedical data analysis. He/she will discover both the fundamentals behind DL techniques and approaches, and their applications on biomedical data. This book can also serve as a reference book for graduate courses in Bioinformatics, AI, ML and DL. The book aims not only at professional researchers and practitioners but also graduate students, senior undergraduate students and young researchers. This book will certainly show the way to new techniques and approaches to make new discoveries.

Explore fundamental to advanced Python 3 topics in six steps, all designed to make you a worthy practitioner. This updated version's approach is based on the "six degrees of separation" theory, which states that everyone and everything is a maximum of six steps away and presents each topic in two parts: theoretical concepts and practical implementation using suitable Python 3 packages. You'll start with the fundamentals of Python 3 programming language, machine learning history, evolution, and the system development frameworks. Key data mining/analysis concepts, such as exploratory analysis, feature dimension reduction, regressions, time series forecasting and their efficient implementation in Scikit-learn are covered as well. You'll also learn commonly used model diagnostic and tuning techniques. These include optimal probability cutoff point for class creation, variance, bias, bagging, boosting, ensemble voting, grid search, random search, Bayesian optimization, and the noise reduction technique for IoT data. Finally, you'll review advanced text mining techniques, recommender systems, neural networks, deep learning, reinforcement learning techniques and their implementation. All the code presented in the book will be available in the form of iPython notebooks to enable you to try out these examples and extend them to your advantage. What You'll Learn Understand machine learning development and frameworks Assess model diagnosis and tuning in machine learning Examine text mining, natural language processing (NLP), and recommender systems Review reinforcement learning and CNN Who This Book Is For Python developers, data engineers, and machine learning engineers looking to expand their knowledge or career into machine learning area.

Discover the practical aspects of implementing deep-learning solutions using the rich Python ecosystem. This book bridges the gap between the academic state-of-the-art and the industry state-of-the-practice by introducing you to deep learning frameworks such as Keras, Theano, and Caffe. The practicalities of these frameworks is often acquired by practitioners by reading source code, manuals, and posting questions on community forums, which tends to be a slow and a painful process. Deep Learning with Python allows you to ramp up to such practical know-how in a short period of time and focus more on the domain, models, and algorithms. This book briefly covers the mathematical prerequisites and fundamentals of deep learning, making this book a good starting point for software developers who want to get started in deep learning. A brief survey of deep learning architectures is also included. Deep Learning with Python also introduces you to key concepts of automatic differentiation and GPU computation which, while not central to deep learning, are critical when it comes to conducting large scale experiments. What You Will Learn Leverage deep learning frameworks in Python namely, Keras, Theano, and Caffe Gain the fundamentals of deep learning with mathematical prerequisites Discover the practical considerations of large scale experiments Take deep learning models to production Who This Book Is For Software developers who want to try out deep learning as a practical solution to a particular problem. Software developers in a data science team who want to take deep learning models developed by data scientists to production.

Train Faster, Reduce Overfitting, and Make Better Predictions
Learning TensorFlow

Deep Learning for Cancer Diagnosis
Mastering Machine Learning with Python in Six Steps

Practical Deep Learning for Cloud, Mobile, and Edge
Theory and Practice of Neural Networks, Computer Vision, Nlp, and Transformers Using Tensorflow

"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
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 module analysis, and where to go next
PART 3 – DEPLOYMENT 15
Deploying to production

This book provides a 'one-stop source' for all readers who are interested in a new, empirical approach to machine learning that, unlike traditional methods, successfully addresses the demands of today's data-driven world. After an introduction to the fundamentals, the book discusses in depth anomaly detection, data partitioning and clustering, as well as classification and predictors. It describes classifiers of zero and first order, and the new, highly efficient and transparent deep rule-based classifiers, particularly highlighting their applications to image processing. Local optimality and stability conditions for the methods presented are formally derived and stated, while the software is also provided as supplemental, open-source material. The book will greatly benefit postgraduate students, researchers and practitioners dealing with advanced data processing, applied mathematicians, software developers of agent-oriented systems, and developers of embedded and real-time systems. It can also be used as a textbook for postgraduate coursework; for this purpose, a standalone set of lecture notes and corresponding lab session notes are available on the same website as the code. Dimitar Filev, Henry Ford Technical Fellow, Ford Motor Company, USA, and Member of the National Academy of Engineering, USA: "The book Empirical Approach to Machine Learning opens new horizons to automated and efficient data processing." Paul J. Werbos, Inventor of the back-propagation method, USA: "I owe great thanks to Professor Plamen Angelov for making this important material available to the community just as I see great practical needs for it, in the new area of making real sense of high-speed data from the brain." Chin-Teng Lin, Distinguished Professor at University of Technology Sydney, Australia: "This new book will set up a milestone for the modern intelligent systems." Edward Tunstel, President of IEEE Systems, Man, Cybernetics Society, USA: "Empirical Approach to Machine Learning provides an insightful and visionary boost of progress in the evolution of computational learning capabilities yielding interpretable and transparent implementations."

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

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.

Better Deep Learning

Dive Into Deep Learning
MACHINE LEARNING

Real-World AI & Computer-Vision Projects Using Python, Keras & TensorFlow
A Guide to Building Deep Learning Systems