
Read Book Pro Deep Learning With TensorFlow A Mathematical Approach To Advanced Artificial Intelligence In Python

Right here, we have countless ebook **Pro Deep Learning With TensorFlow A Mathematical Approach To Advanced Artificial Intelligence In Python** and collections to check out. We additionally come up with the money for variant types and afterward type of the books to browse. The within acceptable limits book, fiction, history, novel, scientific research, as skillfully as various further sorts of books are readily approachable here.

As this Pro Deep Learning With TensorFlow A Mathematical Approach To Advanced Artificial Intelligence In Python, it ends taking place being one of the favored ebook Pro Deep Learning With TensorFlow A Mathematical Approach To Advanced Artificial Intelligence In Python collections that we have. This is why you remain in the best website to look the amazing ebook to have.

KEY=WITH - CARDENAS COHEN

Pro Deep Learning with TensorFlow A Mathematical Approach to Advanced Artificial Intelligence in Python

Apres Deploy deep learning solutions in production with ease using TensorFlow. You'll also develop the mathematical understanding and intuition required to invent new deep learning architectures and solutions on your own. Pro Deep Learning with TensorFlow provides practical, hands-on expertise so you can learn deep learning from scratch and deploy meaningful deep learning solutions. This book will allow you to get up to speed quickly using TensorFlow and to optimize different deep learning architectures. All of the practical aspects of deep learning that are relevant in any industry are emphasized in this book. You will be able to use the prototypes demonstrated to build new deep learning applications. The code presented in the book is available in the form of iPython notebooks and scripts which allow you to try out examples and extend them in interesting ways. You will be equipped with the

mathematical foundation and scientific knowledge to pursue research in this field and give back to the community. What You'll Learn Understand full stack deep learning using TensorFlow and gain a solid mathematical foundation for deep learning Deploy complex deep learning solutions in production using TensorFlow Carry out research on deep learning and perform experiments using TensorFlow Who This Book Is For Data scientists and machine learning professionals, software developers, graduate students, and open source enthusiasts

Pro Deep Learning with TensorFlow

A Mathematical Approach to Advanced Artificial Intelligence in Python

Apress Deploy deep learning solutions in production with ease using TensorFlow. You'll also develop the mathematical understanding and intuition required to invent new deep learning architectures and solutions on your own. Pro Deep Learning with TensorFlow provides practical, hands-on expertise so you can learn deep learning from scratch and deploy meaningful deep learning solutions. This book will allow you to get up to speed quickly using TensorFlow and to optimize different deep learning architectures. All of the practical aspects of deep learning that are relevant in any industry are emphasized in this book. You will be able to use the prototypes demonstrated to build new deep learning applications. The code presented in the book is available in the form of iPython notebooks and scripts which allow you to try out examples and extend them in interesting ways. You will be equipped with the mathematical foundation and scientific knowledge to pursue research in this field and give back to the community. What You'll Learn Understand full stack deep learning using TensorFlow and gain a solid mathematical foundation for deep learning Deploy complex deep learning solutions in production using TensorFlow Carry out research on deep learning and perform experiments using TensorFlow Who This Book Is For Data scientists and machine learning professionals, software developers, graduate students, and open source enthusiasts

Intelligent Projects Using Python

9 real-world AI projects leveraging machine learning and deep learning

with TensorFlow and Keras

Packt Publishing Ltd Implement machine learning and deep learning methodologies to build smart, cognitive AI projects using Python Key Features A go-to guide to help you master AI algorithms and concepts 8 real-world projects tackling different challenges in healthcare, e-commerce, and surveillance Use TensorFlow, Keras, and other Python libraries to implement smart AI applications Book Description This book will be a perfect companion if you want to build insightful projects from leading AI domains using Python. The book covers detailed implementation of projects from all the core disciplines of AI. We start by covering the basics of how to create smart systems using machine learning and deep learning techniques. You will assimilate various neural network architectures such as CNN, RNN, LSTM, to solve critical new world challenges. You will learn to train a model to detect diabetic retinopathy conditions in the human eye and create an intelligent system for performing a video-to-text translation. You will use the transfer learning technique in the healthcare domain and implement style transfer using GANs. Later you will learn to build AI-based recommendation systems, a mobile app for sentiment analysis and a powerful chatbot for carrying customer services. You will implement AI techniques in the cybersecurity domain to generate Captchas. Later you will train and build autonomous vehicles to self-drive using reinforcement learning. You will be using libraries from the Python ecosystem such as TensorFlow, Keras and more to bring the core aspects of machine learning, deep learning, and AI. By the end of this book, you will be skilled to build your own smart models for tackling any kind of AI problems without any hassle. What you will learn Build an intelligent machine translation system using seq-2-seq neural translation machines Create AI applications using GAN and deploy smart mobile apps using TensorFlow Translate videos into text using CNN and RNN Implement smart AI Chatbots, and integrate and extend them in several domains Create smart reinforcement, learning-based applications using Q-Learning Break and generate CAPTCHA using Deep Learning and Adversarial Learning Who this book is for This book is intended for data scientists, machine learning professionals, and deep learning practitioners who are ready to extend their knowledge and potential in AI. If you want to build real-life smart systems to play a crucial role in every complex domain, then this book is what you need. Knowledge of Python programming and a familiarity with basic machine learning and deep learning concepts are expected to help you get the most out of the book

Deep Learning with TensorFlow

Packt Publishing Ltd Delve into neural networks, implement deep learning algorithms, and explore layers of data abstraction with the help of this comprehensive TensorFlow guide About This Book Learn how to implement advanced techniques in deep learning with Google's brainchild, TensorFlow Explore deep neural networks and layers of data abstraction with the help of this comprehensive guide Real-world contextualization through some deep learning problems concerning research and application Who This Book Is For The book is

intended for a general audience of people interested in machine learning and machine intelligence. A rudimentary level of programming in one language is assumed, as is a basic familiarity with computer science techniques and technologies, including a basic awareness of computer hardware and algorithms. Some competence in mathematics is needed to the level of elementary linear algebra and calculus. What You Will Learn Learn about machine learning landscapes along with the historical development and progress of deep learning Learn about deep machine intelligence and GPU computing with the latest TensorFlow 1.x Access public datasets and utilize them using TensorFlow to load, process, and transform data Use TensorFlow on real-world datasets, including images, text, and more Learn how to evaluate the performance of your deep learning models Using deep learning for scalable object detection and mobile computing Train machines quickly to learn from data by exploring reinforcement learning techniques Explore active areas of deep learning research and applications In Detail Deep learning is the step that comes after machine learning, and has more advanced implementations. Machine learning is not just for academics anymore, but is becoming a mainstream practice through wide adoption, and deep learning has taken the front seat. As a data scientist, if you want to explore data abstraction layers, this book will be your guide. This book shows how this can be exploited in the real world with complex raw data using TensorFlow 1.x. Throughout the book, you'll learn how to implement deep learning algorithms for machine learning systems and integrate them into your product offerings, including search, image recognition, and language processing. Additionally, you'll learn how to analyze and improve the performance of deep learning models. This can be done by comparing algorithms against benchmarks, along with machine intelligence, to learn from the information and determine ideal behaviors within a specific context. After finishing the book, you will be familiar with machine learning techniques, in particular the use of TensorFlow for deep learning, and will be ready to apply your knowledge to research or commercial projects. Style and approach This step-by-step guide will explore common, and not so common, deep neural networks and show how these can be exploited in the real world with complex raw data. With the help of practical examples, you will learn how to implement different types of neural nets to build smart applications related to text, speech, and image data processing.

Mathematics for Machine Learning

Cambridge University Press Distills key concepts from linear algebra, geometry, matrices, calculus, optimization, probability and statistics that are used in machine learning.

Deep Learning

A Practitioner's Approach

"O'Reilly Media, Inc." How can machine learning--especially deep neural networks--make a real difference in your organization? This hands-on guide not only provides practical information, but helps you get started building efficient deep learning networks. The authors provide the fundamentals of deep learning--tuning, parallelization, vectorization, and building pipelines--that are valid for any library before introducing the 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.

Quantum Machine Learning With Python

Using Cirq from Google Research and IBM Qiskit

Apress Quickly scale up to Quantum computing and Quantum machine learning foundations and related mathematics and expose them to different use cases that can be solved through Quantum based algorithms. This book explains Quantum Computing, which leverages the Quantum mechanical properties sub-atomic particles. It also examines Quantum machine learning, which can help solve some of the most challenging problems in forecasting, financial modeling, genomics, cybersecurity, supply chain logistics, cryptography among others. You'll start by reviewing the fundamental concepts of Quantum Computing, such as Dirac Notations, Qubits, and Bell state, followed by postulates and mathematical foundations of Quantum Computing. Once the foundation base is set, you'll delve deep into Quantum based algorithms including Quantum Fourier transform, phase estimation, and HHL (Harrow-Hassidim-Lloyd) among others. You'll then be introduced to Quantum machine learning and Quantum deep learning-based algorithms, along with advanced topics of Quantum adiabatic processes and Quantum based optimization. Throughout the book, there are Python implementations of different Quantum machine learning and Quantum computing algorithms using the Qiskit toolkit from IBM and Cirq from Google Research. What You'll Learn Understand Quantum computing and Quantum machine learning Explore varied domains and the scenarios where Quantum machine learning solutions can be applied Develop expertise in algorithm development in varied Quantum computing frameworks Review the major challenges of building large scale Quantum computers and applying its various techniques Who This Book Is For Machine Learning enthusiasts and engineers who want to quickly scale up to Quantum Machine Learning

Learning TensorFlow

A Guide to Building Deep Learning Systems

"O'Reilly Media, Inc." 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

Mathematics Rebooted

A Fresh Approach to Understanding

Oxford University Press This book is about mathematical thinking, learning and understanding. It is about ways in which good representations capture mathematical ideas, and about building broad and deep knowledge by understanding the links between those ideas. It draws on research in mathematics education and psychology to explain why some misunderstandings and confusions arise for almost everyone, and it describes ways to think about mathematical ideas correctly and with confidence.

Deep Learning with Python

Simon and Schuster 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-unusable speech and image recognition, to near-human accuracy. We went from machines that couldn't beat a serious Go player, 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 Learning Representations (ICLR), and others. Table of Contents PART 1 - 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 for Natural Language Processing

Creating Neural Networks with Python

Apress Discover the concepts of deep learning used for natural language processing (NLP), with full-fledged examples of neural network models such as recurrent neural networks, long short-term memory networks, and sequence-2-sequence models.

You'll start by covering the mathematical prerequisites and the fundamentals of deep learning and NLP with practical examples. The first three chapters of the book cover the basics of NLP, starting with word-vector representation before moving onto advanced algorithms. The final chapters focus entirely on implementation, and deal with sophisticated architectures such as RNN, LSTM, and Seq2seq, using Python tools: TensorFlow, and Keras. Deep Learning for Natural Language Processing follows a progressive approach and combines all the knowledge you have gained to build a question-answer chatbot system. This book is a good starting point for people who want to get started in deep learning for NLP. All the code presented in the book will be available in the form of IPython notebooks and scripts, which allow you to try out the examples and extend them in interesting ways. What You Will Learn Gain the fundamentals of deep learning and its mathematical prerequisites Discover deep learning frameworks in Python Develop a chatbot Implement a research paper on sentiment classification Who This Book Is For Software developers who are curious to try out deep learning with NLP.

Machine Learning with TensorFlow

Manning Publications Summary Machine Learning with TensorFlow gives readers a solid foundation in machine-learning concepts plus hands-on experience coding TensorFlow with Python. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology TensorFlow, Google's library for large-scale machine learning, simplifies often-complex computations by representing them as graphs and efficiently mapping parts of the graphs to machines in a cluster or to the processors of a single machine. About the Book Machine Learning with TensorFlow gives readers a solid foundation in machine-learning concepts plus hands-on experience coding TensorFlow with Python. You'll learn the basics by working with classic prediction, classification, and clustering algorithms. Then, you'll move on to the money chapters: exploration of deep-learning concepts like autoencoders, recurrent neural networks, and reinforcement learning. Digest this book and you will be ready to use TensorFlow for machine-learning and deep-learning applications of your own. What's Inside Matching your tasks to the right machine-learning and deep-learning approaches Visualizing algorithms with TensorBoard Understanding and using neural networks About the Reader Written for developers experienced with Python and algebraic concepts like vectors and matrices. About the Author Author Nishant Shukla is a computer vision researcher focused on applying machine-learning techniques in robotics. Senior technical editor, Kenneth Fricklas, is a seasoned developer, author, and machine-learning practitioner. Table of Contents PART 1 - YOUR MACHINE-LEARNING RIG A machine-learning odyssey TensorFlow essentials PART 2 - CORE LEARNING ALGORITHMS Linear regression and beyond A gentle introduction to classification Automatically clustering data Hidden Markov models PART 3 - THE NEURAL NETWORK PARADIGM A peek into autoencoders Reinforcement learning Convolutional neural networks Recurrent neural networks Sequence-to-sequence models for chatbots Utility landscape

Deep Learning with JavaScript

Neural networks in TensorFlow.js

Simon and Schuster Summary Deep learning has transformed the fields of computer vision, image processing, and natural language applications. Thanks to TensorFlow.js, now JavaScript developers can build deep learning apps without relying on Python or R. Deep Learning with JavaScript shows developers how they can bring DL technology to the web. Written by the main authors of the TensorFlow library, this new book provides fascinating use cases and in-depth instruction for deep learning apps in JavaScript in your browser or on Node. Foreword by Nikhil Thorat and Daniel Smilkov. About the technology Running deep learning applications in the browser or on Node-based backends opens up exciting possibilities for smart web applications. With the TensorFlow.js library, you build and train deep learning models with JavaScript. Offering uncompromising production-quality scalability, modularity, and responsiveness, TensorFlow.js really shines for its portability. Its models run anywhere JavaScript runs, pushing ML farther up the application stack. About the book In Deep Learning with JavaScript, you'll learn to use TensorFlow.js to build deep learning models that run directly in the browser. This fast-paced book, written by Google engineers, is practical, engaging, and easy to follow. Through diverse examples featuring text analysis, speech processing, image recognition, and self-learning game AI, you'll master all the basics of deep learning and explore advanced concepts, like retraining existing models for transfer learning and image generation. What's inside - Image and language processing in the browser - Tuning ML models with client-side data - Text and image creation with generative deep learning - Source code samples to test and modify About the reader For JavaScript programmers interested in deep learning. About the author Shanging Cai, Stanley Bileschi and Eric D. Nielsen are software engineers with experience on the Google Brain team, and were crucial to the development of the high-level API of TensorFlow.js. This book is based in part on the classic, Deep Learning with Python by François Chollet. TOC: PART 1 - MOTIVATION AND BASIC CONCEPTS 1 • Deep learning and JavaScript PART 2 - A GENTLE INTRODUCTION TO TENSORFLOW.JS 2 • Getting started: Simple linear regression in TensorFlow.js 3 • Adding nonlinearity: Beyond weighted sums 4 • Recognizing images and sounds using convnets 5 • Transfer learning: Reusing pretrained neural networks PART 3 - ADVANCED DEEP LEARNING WITH TENSORFLOW.JS 6 • Working with data 7 • Visualizing data and models 8 • Underfitting, overfitting, and the universal workflow of machine learning 9 • Deep learning for sequences and text 10 • Generative deep learning 11 • Basics of deep reinforcement learning PART 4 - SUMMARY AND CLOSING WORDS 12 • Testing, optimizing, and deploying models 13 • Summary, conclusions, and beyond

Practical Machine Learning and Image Processing

For Facial Recognition, Object Detection, and Pattern Recognition Using Python

Apres Gain insights into image-processing methodologies and algorithms, using machine learning and neural networks in Python. This book begins with the environment setup, understanding basic image-processing terminology, and exploring Python concepts that will be useful for implementing the algorithms discussed in the book. You will then cover all the core image processing algorithms in detail before moving onto the biggest computer vision library: OpenCV. You'll see the OpenCV algorithms and how to use them for image processing. The next section looks at advanced machine learning and deep learning methods for image processing and classification. You'll work with concepts such as pulse coupled neural networks, AdaBoost, XG boost, and convolutional neural networks for image-specific applications. Later you'll explore how models are made in real time and then deployed using various DevOps tools. All the concepts in Practical Machine Learning and Image Processing are explained using real-life scenarios. After reading this book you will be able to apply image processing techniques and make machine learning models for customized application. What You Will Learn Discover image-processing algorithms and their applications using Python Explore image processing using the OpenCV library Use TensorFlow, scikit-learn, NumPy, and other libraries Work with machine learning and deep learning algorithms for image processing Apply image-processing techniques to five real-time projects Who This Book Is For Data scientists and software developers interested in image processing and computer vision.

Deep Learning with PyTorch

Manning Publications 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, providing a comfortable Python experience that gets you started quickly and then grows with you as you—and your deep learning skills—become more sophisticated. Deep Learning with PyTorch will make that journey engaging and fun. Summary 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, providing a comfortable Python experience that gets you started quickly and then

grows with you as you—and your deep learning skills—become more sophisticated. *Deep Learning with PyTorch* will make that journey engaging and fun. Foreword by Soumith Chintala, Cocreator of PyTorch. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the technology Although many deep learning tools use Python, the PyTorch library is truly Pythonic. Instantly familiar to anyone who knows PyData tools like NumPy and scikit-learn, PyTorch simplifies deep learning without sacrificing advanced features. It's excellent for building quick models, and it scales smoothly from laptop to enterprise. Because companies like Apple, Facebook, and JPMorgan Chase rely on PyTorch, it's a great skill to have as you expand your career options. It's easy to get started with PyTorch. It minimizes cognitive overhead without sacrificing the access to advanced features, meaning you can focus on what matters the most - building and training the latest and greatest deep learning models and contribute to making a dent in the world. PyTorch is also a snap to scale and extend, and it partners well with other Python tooling. PyTorch has been adopted by hundreds of deep learning practitioners and several first-class players like FAIR, OpenAI, FastAI and Purdue. About the book *Deep Learning with PyTorch* teaches you to create neural networks and deep learning systems with PyTorch. This practical book quickly gets you to work building a real-world example from scratch: a tumor image classifier. Along the way, it covers best practices for the entire DL pipeline, including the PyTorch Tensor API, loading data in Python, monitoring training, and visualizing results. After covering the basics, the book will take you on a journey through larger projects. The centerpiece of the book is a neural network designed for cancer detection. You'll discover ways for training networks with limited inputs and start processing data to get some results. You'll sift through the unreliable initial results and focus on how to diagnose and fix the problems in your neural network. Finally, you'll look at ways to improve your results by training with augmented data, make improvements to the model architecture, and perform other fine tuning. What's inside

Training deep neural networks
Implementing modules and loss functions
Utilizing pretrained models from PyTorch Hub
Exploring code samples in Jupyter Notebooks

About the reader For Python programmers with an interest in machine learning. About the author Eli Stevens had roles from software engineer to CTO, and is currently working on machine learning in the self-driving-car industry. Luca Antiga is cofounder of an AI engineering company and an AI tech startup, as well as a former PyTorch contributor. Thomas Viehmann is a PyTorch core developer and machine learning trainer and consultant. 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

Short-Term Load Forecasting 2019

MDPI Short-term load forecasting (STLF) plays a key role in the formulation of economic, reliable, and secure operating strategies (planning, scheduling, maintenance, and control processes, among others) for a power system and will be significant in the future. However, there is still much to do in these research areas. The deployment of enabling technologies (e.g., smart meters) has made high-granularity data available for many customer segments and to approach many issues, for instance, to make forecasting tasks feasible at several demand aggregation levels. The first challenge is the improvement of STLF models and their performance at new aggregation levels. Moreover, the mix of renewables in the power system, and the necessity to include more flexibility through demand response initiatives have introduced greater uncertainties, which means new challenges for STLF in a more dynamic power system in the 2030-50 horizon. Many techniques have been proposed and applied for STLF, including traditional statistical models and AI techniques. Besides, distribution planning needs, as well as grid modernization, have initiated the development of hierarchical load forecasting. Analogously, the need to face new sources of uncertainty in the power system is giving more importance to probabilistic load forecasting. This Special Issue deals with both fundamental research and practical application research on STLF methodologies to face the challenges of a more distributed and customer-centered power system.

Math for Deep Learning

What You Need to Know to Understand Neural Networks

No Starch Press Math for Deep Learning provides the essential math you need to understand deep learning discussions, explore more complex implementations, and better use the deep learning toolkits. With Math for Deep Learning, you'll learn the essential mathematics used by and as a background for deep learning. You'll work through Python examples to learn key deep learning related topics in probability, statistics, linear algebra, differential calculus, and matrix calculus as well as how to implement data flow in a neural network, backpropagation, and gradient descent. You'll also use Python to work through the mathematics that underlies those algorithms and even build a fully-functional neural network. In addition you'll find coverage of gradient descent including variations commonly used by the deep learning community: SGD, Adam, RMSprop, and Adagrad/Adadelta.

Introduction to Deep Learning

From Logical Calculus to Artificial Intelligence

Springer This textbook presents a concise, accessible and engaging first introduction to deep learning, offering a wide range of connectionist models which represent the current state-of-the-art. The text explores the most popular algorithms and architectures in a simple and intuitive style, explaining the mathematical derivations in a step-by-step manner. The content coverage includes convolutional networks, LSTMs, Word2vec, RBMs, DBNs, neural Turing machines, memory networks and autoencoders. Numerous examples in working Python code are provided throughout the book, and the code is also supplied separately at an accompanying website. Topics and features: introduces the fundamentals of machine learning, and the mathematical and computational prerequisites for deep learning; discusses feed-forward neural networks, and explores the modifications to these which can be applied to any neural network; examines convolutional neural networks, and the recurrent connections to a feed-forward neural network; describes the notion of distributed representations, the concept of the autoencoder, and the ideas behind language processing with deep learning; presents a brief history of artificial intelligence and neural networks, and reviews interesting open research problems in deep learning and connectionism. This clearly written and lively primer on deep learning is essential reading for graduate and advanced undergraduate students of computer science, cognitive science and mathematics, as well as fields such as linguistics, logic, philosophy, and psychology.

Machine Learning with TensorFlow, Second Edition

Manning Publications Updated with new code, new projects, and new chapters, *Machine Learning with TensorFlow, Second Edition* gives readers a solid foundation in machine-learning concepts and the TensorFlow library. Summary Updated with new code, new projects, and new chapters, *Machine Learning with TensorFlow, Second Edition* gives readers a solid foundation in machine-learning concepts and the TensorFlow library. Written by NASA JPL Deputy CTO and Principal Data Scientist Chris Mattmann, all examples are accompanied by downloadable Jupyter Notebooks for a hands-on experience coding TensorFlow with Python. New and revised content expands coverage of core machine learning algorithms, and advancements in neural networks such as VGG-Face facial identification classifiers and deep speech classifiers. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the technology Supercharge your data analysis with machine learning! ML algorithms automatically improve as they

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

Deep Learning for Coders with fastai and PyTorch

O'Reilly Media *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 consistent interface to the 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 in computer vision, 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 Consider the ethical implications of your work Gain

insight from the foreword by PyTorch cofounder, Soumith Chintala

Progress in Advanced Computing and Intelligent Engineering

Proceedings of ICACIE 2020

Springer Nature This book focuses on theory, practice and applications in the broad areas of advanced computing techniques and intelligent engineering. This book includes 74 scholarly articles which were accepted for presentation from 294 submissions in the 5th ICACIE during 25-27 June 2020 at Université des Mascareignes (UdM), Mauritius, in collaboration with Rama Devi Women's University, Bhubaneswar, India, and S'O'A Deemed to be University, Bhubaneswar, India. This book brings together academicians, industry persons, research scholars and students to share and disseminate their knowledge and scientific research work related to advanced computing and intelligent engineering. It helps to provide a platform to the young researchers to find the practical challenges encountered in these areas of research and the solutions adopted. The book helps to disseminate the knowledge about some innovative and active research directions in the field of advanced computing techniques and intelligent engineering, along with some current issues and applications of related topics.

Digital Health in Focus of Predictive, Preventive and Personalised Medicine

Springer Nature The edition will cover proceedings of the second International conference on digital health Technologies (ICDHT 2019). The conference will address the topic of P4 medicine from the information technology point of view, and will be focused on the following topics: - Artificial Intelligence for health • Knowledge extraction • Decision-aid systems • Data analysis and risk prediction • Machine learning, deep learning - Health data processing • Data preprocessing, cleaning, management and mining • Computer-aided detection • Big data analysis, prediction and prevention • Cognitive algorithms for healthcare handling dynamic context management • Augmented reality, Motion detection and activity recognition - Devices, infrastructure and communication • Wearable & connected devices • Communication infrastructures, architectures and standards Blockchain for e-Health • Computing/storage infrastructures for e-Health • IoT devices & architectures for Smart Healthcare - Health information systems • Telemedicine, Teleservices • Computing/storage infrastructures for e-Health • Clinical Data Visualisation Standards - Security and privacy for e-health • Health data Analytics for Security and

Privacy • E-health Software and Hardware Security • Embedded Security for e-health
- Applications in P4 medicine

Avatar-Based Control, Estimation, Communications, and Development of Neuron Multi-Functional Technology Platforms

IGI Global Competition in today's global market offers strong motivation for the development of sophisticated tools within computer science. The neuron multi-functional technology platform is a developing field of study that regards the various interactive approaches that can be applied within this subject matter. As advancing technologies continue to emerge, managers and researchers need a compilation of research that discusses the advancements and specific implementations of these intelligent approaches with this platform. Avatar-Based Control, Estimation, Communications, and Development of Neuron Multi-Functional Technology Platforms is a pivotal reference source that provides vital research on the application of artificial and natural approaches towards neuron-based programs. While highlighting topics such as natural intelligence, neurolinguistics, and smart data storage, this publication presents techniques, case studies, and methodologies that combine the use of intelligent artificial and natural approaches with optimization techniques for facing problems and combines many types of hardware and software with a variety of communication technologies to enable the development of innovative applications. This book is ideally designed for researchers, practitioners, scientists, field experts, professors, and students seeking current research on the optimization of avatar-based advancements in multifaceted technology systems.

Applied Natural Language Processing with Python

Implementing Machine Learning and Deep Learning Algorithms for Natural Language Processing

Apress Learn to harness the power of AI for natural language processing, performing tasks such as spell check, text summarization, document classification, and natural language generation. Along the way, you will learn the skills to implement these

methods in larger infrastructures to replace existing code or create new algorithms. *Applied Natural Language Processing with Python* starts with reviewing the necessary machine learning concepts before moving onto discussing various NLP problems. After reading this book, you will have the skills to apply these concepts in your own professional environment. *What You Will Learn* Utilize various machine learning and natural language processing libraries such as TensorFlow, Keras, NLTK, and Gensim Manipulate and preprocess raw text data in formats such as .txt and .pdf Strengthen your skills in data science by learning both the theory and the application of various algorithms *Who This Book Is For* You should be at least a beginner in ML to get the most out of this text, but you needn't feel that you need be an expert to understand the content.

Dive Into Deep Learning Tools for Engagement

Corwin Press 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.

Hands-On Machine Learning with R

CRC Press *Hands-on Machine Learning with R* provides a practical and applied approach to learning and developing intuition into today's most popular machine learning methods. This book serves as a practitioner's guide to the machine learning process and is meant to help the reader learn to apply the machine learning stack

within R, which includes using various R packages such as *glmnet*, *h2o*, *ranger*, *xgboost*, *keras*, and others to effectively model and gain insight from their data. The book favors a hands-on approach, providing an intuitive understanding of machine learning concepts through concrete examples and just a little bit of theory. Throughout this book, the reader will be exposed to the entire machine learning process including feature engineering, resampling, hyperparameter tuning, model evaluation, and interpretation. The reader will be exposed to powerful algorithms such as regularized regression, random forests, gradient boosting machines, deep learning, generalized low rank models, and more! By favoring a hands-on approach and using real world data, the reader will gain an intuitive understanding of the architectures and engines that drive these algorithms and packages, understand when and how to tune the various hyperparameters, and be able to interpret model results. By the end of this book, the reader should have a firm grasp of R's machine learning stack and be able to implement a systematic approach for producing high quality modeling results. Features:

- Offers a practical and applied introduction to the most popular machine learning methods.
- Topics covered include feature engineering, resampling, deep learning and more.
- Uses a hands-on approach and real world data.

Learn Keras for Deep Neural Networks

A Fast-Track Approach to Modern Deep Learning with Python

Apres Learn, understand, and implement deep neural networks in a math- and programming-friendly approach using Keras and Python. The book focuses on an end-to-end approach to developing supervised learning algorithms in regression and classification with practical business-centric use-cases implemented in Keras. The overall book comprises three sections with two chapters in each section. The first section prepares you with all the necessary basics to get started in deep learning. Chapter 1 introduces you to the world of deep learning and its difference from machine learning, the choices of frameworks for deep learning, and the Keras ecosystem. You will cover a real-life business problem that can be solved by supervised learning algorithms with deep neural networks. You'll tackle one use case for regression and another for classification leveraging popular Kaggle datasets. Later, you will see an interesting and challenging part of deep learning: hyperparameter tuning; helping you further improve your models when building robust deep learning applications. Finally, you'll further hone your skills in deep learning and cover areas of active development and research in deep learning. At the end of Learn Keras for Deep Neural Networks, you will have a thorough understanding of deep learning principles and have practical hands-on experience in

developing enterprise-grade deep learning solutions in Keras. What You'll Learn Master fast-paced practical deep learning concepts with math- and programming-friendly abstractions. Design, develop, train, validate, and deploy deep neural networks using the Keras framework Use best practices for debugging and validating deep learning models Deploy and integrate deep learning as a service into a larger software service or product Extend deep learning principles into other popular frameworks Who This Book Is For Software engineers and data engineers with basic programming skills in any language and who are keen on exploring deep learning for a career move or an enterprise project.

TensorFlow For Dummies

John Wiley & Sons Become a machine learning pro! Google TensorFlow has become the darling of financial firms and research organizations, but the technology can be intimidating and the learning curve is steep. Luckily, TensorFlow For Dummies is here to offer you a friendly, easy-to-follow book on the subject. Inside, you'll find out how to write applications with TensorFlow, while also grasping the concepts underlying machine learning—all without ever losing your cool! Machine learning has become ubiquitous in modern society, and its applications include language translation, robotics, handwriting analysis, financial prediction, and image recognition. TensorFlow is Google's preeminent toolset for machine learning, and this hands-on guide makes it easy to understand, even for those without a background in artificial intelligence. Install TensorFlow on your computer Learn the fundamentals of statistical regression and neural networks Visualize the machine learning process with TensorBoard Perform image recognition with convolutional neural networks (CNNs) Analyze sequential data with recurrent neural networks (RNNs) Execute TensorFlow on mobile devices and the Google Cloud Platform (GCP) If you're a manager or software developer looking to use TensorFlow for machine learning, this is the book you'll want to have close by.

Deep Learning with R

Simon and Schuster Summary Deep Learning with R introduces the world of deep learning using the powerful Keras library and its R language interface. The book builds your understanding of deep learning through intuitive explanations and practical examples. Continue your journey into the world of deep learning with Deep Learning with R in Motion, a practical, hands-on video course available exclusively at Manning.com (www.manning.com/livevideo/deep-learning-with-r-in-motion). 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. Deep-learning systems now enable previously impossible smart applications, revolutionizing image recognition and natural-language processing, and identifying complex patterns in data. The Keras deep-learning library provides data scientists and developers working in R a state-of-the-art toolset for tackling deep-learning tasks. About the Book Deep Learning with R introduces the world of deep learning using the powerful Keras library and its R

language interface. Initially written for Python as *Deep Learning with Python* by Keras creator and Google AI researcher François Chollet and adapted for R by RStudio founder J. J. Allaire, this book builds your understanding of deep learning through intuitive explanations and practical examples. You'll practice your new skills with R-based applications in computer vision, natural-language processing, and generative models. What's Inside Deep learning from first principles Setting up your own deep-learning environment Image classification and generation Deep learning for text and sequences About the Reader You'll need intermediate R programming skills. No previous experience with machine learning or deep learning is assumed. About the Authors François Chollet is a deep-learning researcher at Google and the author of the Keras library. J.J. Allaire is the founder of RStudio and the author of the R interfaces to TensorFlow and Keras. Table of Contents PART 1 - 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

Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow

Concepts, Tools, and Techniques to Build Intelligent Systems

O'Reilly Media Through a series of recent breakthroughs, deep learning has boosted the entire field of machine learning. Now, even programmers who know close to nothing about this technology can use simple, efficient tools to implement programs capable of learning from data. This practical book shows you how. By using concrete examples, minimal theory, and two production-ready Python frameworks—Scikit-Learn and TensorFlow—author Aurélien Géron helps you gain an intuitive understanding of the concepts and tools for building intelligent systems. You'll learn a range of techniques, starting with simple linear regression and progressing to deep neural networks. With exercises in each chapter to help you apply what you've learned, all you need is programming experience to get started. Explore the machine learning landscape, particularly neural nets Use Scikit-Learn to track an example machine-learning project end-to-end Explore several training models, including support vector machines, decision trees, random forests, and ensemble methods Use the TensorFlow library to build and train neural nets Dive into neural net architectures, including convolutional nets, recurrent nets, and deep reinforcement learning Learn techniques for training and scaling deep neural nets

Bayesian Methods for Hackers

Probabilistic Programming and Bayesian Inference

Addison-Wesley Professional Master Bayesian Inference through Practical Examples and Computation—Without Advanced Mathematical Analysis Bayesian methods of inference are deeply natural and extremely powerful. However, most discussions of Bayesian inference rely on intensely complex mathematical analyses and artificial examples, making it inaccessible to anyone without a strong mathematical background. Now, though, Cameron Davidson-Pilon introduces Bayesian inference from a computational perspective, bridging theory to practice—freeing you to get results using computing power. Bayesian Methods for Hackers illuminates Bayesian inference through probabilistic programming with the powerful PyMC language and the closely related Python tools NumPy, SciPy, and Matplotlib. Using this approach, you can reach effective solutions in small increments, without extensive mathematical intervention. Davidson-Pilon begins by introducing the concepts underlying Bayesian inference, comparing it with other techniques and guiding you through building and training your first Bayesian model. Next, he introduces PyMC through a series of detailed examples and intuitive explanations that have been refined after extensive user feedback. You'll learn how to use the Markov Chain Monte Carlo algorithm, choose appropriate sample sizes and priors, work with loss functions, and apply Bayesian inference in domains ranging from finance to marketing. Once you've mastered these techniques, you'll constantly turn to this guide for the working PyMC code you need to jumpstart future projects. Coverage includes

- Learning the Bayesian “state of mind” and its practical implications
- Understanding how computers perform Bayesian inference
- Using the PyMC Python library to program Bayesian analyses
- Building and debugging models with PyMC
- Testing your model’s “goodness of fit”
- Opening the “black box” of the Markov Chain Monte Carlo algorithm to see how and why it works
- Leveraging the power of the “Law of Large Numbers”
- Mastering key concepts, such as clustering, convergence, autocorrelation, and thinning
- Using loss functions to measure an estimate’s weaknesses based on your goals and desired outcomes
- Selecting appropriate priors and understanding how their influence changes with dataset size
- Overcoming the “exploration versus exploitation” dilemma: deciding when “pretty good” is good enough
- Using Bayesian inference to improve A/B testing
- Solving data science problems when only small amounts of data are available

Cameron Davidson-Pilon has worked in many areas of applied mathematics, from the evolutionary dynamics of genes and diseases to stochastic modeling of financial prices. His contributions to the open source community include lifelines, an implementation of survival analysis in Python. Educated at the University of Waterloo and at the Independent University of Moscow, he currently works with the online commerce leader Shopify.

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

Addison-Wesley Professional 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 statistic 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 research (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

Deep Learning on Graphs

Cambridge University Press A comprehensive text on foundations and techniques of graph neural networks with applications in NLP, data mining, vision and healthcare.

Interpretable Machine Learning

Lulu.com

Probability for Statistics and Machine Learning

Fundamentals and Advanced Topics

Springer Science & Business Media This book provides a versatile and lucid treatment of classic as well as modern probability theory, while integrating them with core topics in statistical theory and also some key tools in machine learning. It is written in an extremely accessible style, with elaborate motivating discussions and numerous worked out examples and exercises. The book has 20 chapters on a wide range of topics, 423 worked out examples, and 808 exercises. It is unique in its unification of probability and statistics, its coverage and its superb exercise sets, detailed bibliography, and in its substantive treatment of many topics of current importance. This book can be used as a text for a year long graduate course in statistics, computer science, or mathematics, for self-study, and as an invaluable research reference on probability and its applications. Particularly worth mentioning are the treatments of distribution theory, asymptotics, simulation and Markov Chain Monte Carlo, Markov chains and martingales, Gaussian processes, VC theory, probability metrics, large deviations, bootstrap, the EM algorithm, confidence intervals, maximum likelihood and Bayes estimates, exponential families, kernels, and Hilbert spaces, and a self contained complete review of univariate probability.

Practical Deep Learning for Cloud, Mobile, and Edge

Real-World AI & Computer-Vision Projects Using Python, Keras & TensorFlow

O'Reilly Media 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

AI for Data Science

Artificial Intelligence Frameworks and Functionality for Deep Learning, Optimization, and Beyond

Technics Publications Master the approaches and principles of Artificial Intelligence (AI) algorithms, and apply them to Data Science projects with Python and Julia code. Aspiring and practicing Data Science and AI professionals, along with Python and Julia programmers, will practice numerous AI algorithms and develop a more holistic understanding of the field of AI, and will learn when to use each framework to tackle projects in our increasingly complex world. The first two chapters introduce the field, with Chapter 1 surveying Deep Learning models and Chapter 2 providing an overview of algorithms beyond Deep Learning, including Optimization, Fuzzy Logic, and Artificial Creativity. The next chapters focus on AI frameworks; they contain data and Python and Julia code in a provided Docker, so you can practice. Chapter 3 covers Apache's MXNet, Chapter 4 covers TensorFlow, and Chapter 5 investigates Keras. After covering these Deep Learning frameworks, we explore a series of optimization frameworks, with Chapter 6 covering Particle Swarm Optimization (PSO), Chapter 7 on Genetic Algorithms (GAs), and Chapter 8 discussing Simulated Annealing (SA). Chapter 9 begins our exploration of advanced AI methods, by covering Convolutional Neural Networks (CNNs) and Recurrent Neural Networks (RNNs). Chapter 10 discusses optimization ensembles and how they can add value to the Data Science pipeline. Chapter 11 contains several alternative AI frameworks including Extreme Learning Machines (ELMs), Capsule Networks (CapsNets), and Fuzzy Inference Systems (FIS). Chapter 12 covers other considerations complementary to the AI topics covered, including Big Data concepts, Data Science specialization areas, and useful data resources to experiment on. A comprehensive glossary is included, as well as a series of appendices covering Transfer Learning, Reinforcement Learning, Autoencoder Systems, and Generative Adversarial Networks. There is also an appendix on the business aspects of AI in data science projects, and an appendix on how to use the Docker image to access the book's data and code. The field of AI is vast, and can be overwhelming for the newcomer to approach. This book will arm you with a solid understanding of the field, plus inspire you to explore further.

AI Crash Course

A fun and hands-on introduction to machine learning, reinforcement learning, deep learning, and artificial intelligence with Python

Packt Publishing Ltd Unlock the power of artificial intelligence with top Udemy AI instructor Hadelin de Ponteves. Key Features Learn from friendly, plain English explanations and practical activities Put ideas into action with 5 hands-on projects that show step-by-step how to build intelligent software Use AI to win classic video games and construct a virtual self-driving car Book Description Welcome to the Robot World ... and start building intelligent software now! Through his best-selling video courses, Hadelin de Ponteves has taught hundreds of thousands of people to write AI software. Now, for the first time, his hands-on, energetic approach is available as a book. Starting with the basics before easing you into more complicated formulas and notation, AI Crash Course gives you everything you need to build AI systems with reinforcement learning and deep learning. Five full working projects put the ideas into action, showing step-by-step how to build intelligent software using the best and easiest tools for AI programming, including Python, TensorFlow, Keras, and PyTorch. AI Crash Course teaches everyone to build an AI to work in their applications. Once you've read this book, you're only limited by your imagination. What you will learn Master the basics of AI without any previous experience Build fun projects, including a virtual-self-driving car and a robot warehouse worker Use AI to solve real-world business problems Learn how to code in Python Discover the 5 principles of reinforcement learning Create your own AI toolkit Who this book is for If you want to add AI to your skillset, this book is for you. It doesn't require data science or machine learning knowledge. Just maths basics (high school level).

Machine Learning in Action

Simon and Schuster Summary Machine Learning in Action is unique book that blends the foundational theories of machine learning with the practical realities of building tools for everyday data analysis. You'll use the flexible Python programming language to build programs that implement algorithms for data classification, forecasting, recommendations, and higher-level features like summarization and simplification. About the Book A machine is said to learn when its performance improves with experience. Learning requires algorithms and programs that capture data and ferret out the interesting or useful patterns. Once the specialized domain of analysts and mathematicians, machine learning is becoming a skill needed by many.

Machine Learning in Action is a clearly written tutorial for developers. It avoids academic language and takes you straight to the techniques you'll use in your day-to-day work. Many (Python) examples present the core algorithms of statistical data processing, data analysis, and data visualization in code you can reuse. You'll understand the concepts and how they fit in with tactical tasks like classification, forecasting, recommendations, and higher-level features like summarization and simplification. Readers need no prior experience with machine learning or statistical processing. Familiarity with Python is helpful. Purchase of the print book comes with an offer of a free PDF, ePub, and Kindle eBook from Manning. Also available is all code from the book. What's Inside A no-nonsense introduction Examples showing common ML tasks Everyday data analysis Implementing classic algorithms like Apriori and Adaboos Table of Contents PART 1 CLASSIFICATION Machine learning basics Classifying with k-Nearest Neighbors Splitting datasets one feature at a time: decision trees Classifying with probability theory: naïve Bayes Logistic regression Support vector machines Improving classification with the AdaBoost meta algorithm PART 2 FORECASTING NUMERIC VALUES WITH REGRESSION Predicting numeric values: regression Tree-based regression PART 3 UNSUPERVISED LEARNING Grouping unlabeled items using k-means clustering Association analysis with the Apriori algorithm Efficiently finding frequent itemsets with FP-growth PART 4 ADDITIONAL TOOLS Using principal component analysis to simplify data Simplifying data with the singular value decomposition Big data and MapReduce

Deep Learning with TensorFlow 2 and Keras

Regression, ConvNets, GANs, RNNs, NLP, and more with TensorFlow 2 and the Keras API, 2nd Edition

Packt Publishing Ltd Build machine and deep learning systems with the newly released TensorFlow 2 and Keras for the lab, production, and mobile devices Key FeaturesIntroduces and then uses TensorFlow 2 and Keras right from the startTeaches key machine and deep learning techniquesUnderstand the fundamentals of deep learning and machine learning through clear explanations and extensive code samplesBook Description Deep Learning with TensorFlow 2 and Keras, Second Edition teaches neural networks and deep learning techniques alongside TensorFlow (TF) and Keras. You'll learn how to write deep learning applications in the most powerful, popular, and scalable machine learning stack available. TensorFlow is the machine learning library of choice for professional applications, while Keras offers a simple and powerful Python API for accessing TensorFlow. TensorFlow 2 provides full Keras integration, making advanced machine

learning easier and more convenient than ever before. This book also introduces neural networks with TensorFlow, runs through the main applications (regression, ConvNets (CNNs), GANs, RNNs, NLP), covers two working example apps, and then dives into TF in production, TF mobile, and using TensorFlow with AutoML. What you will learn

Build machine learning and deep learning systems with TensorFlow 2 and the Keras API

Use Regression analysis, the most popular approach to machine learning

Understand ConvNets (convolutional neural networks) and how they are essential for deep learning systems such as image classifiers

Use GANs (generative adversarial networks) to create new data that fits with existing patterns

Discover RNNs (recurrent neural networks) that can process sequences of input intelligently, using one part of a sequence to correctly interpret another

Apply deep learning to natural human language and interpret natural language texts to produce an appropriate response

Train your models on the cloud and put TF to work in real environments

Explore how Google tools can automate simple ML workflows without the need for complex modeling

Who this book is for This book is for Python developers and data scientists who want to build machine learning and deep learning systems with TensorFlow. This book gives you the theory and practice required to use Keras, TensorFlow 2, and AutoML to build machine learning systems. Some knowledge of machine learning is expected.