

Deep Learning Deep Learning Explained To Your Granny A Guide For Beginners Machine Learning

Right here, we have countless ebook **Deep Learning Deep Learning Explained To Your Granny A Guide For Beginners Machine Learning** and collections to check out. We additionally pay for variant types and next type of the books to browse. The all right book, fiction, history, novel, scientific research, as competently as various other sorts of books are readily clear here.

As this Deep Learning Deep Learning Explained To Your Granny A Guide For Beginners Machine Learning , it ends occurring innate one of the favored books Deep Learning Deep Learning Explained To Your Granny A Guide For Beginners Machine Learning collections that we have. This is why you remain in the best website to see the amazing books to have.

Explainable AI: Interpreting, Explaining and Visualizing Deep Learning - Wojciech Samek
2019-09-10

The development of “intelligent” systems that can take decisions and perform autonomously might lead to faster and more consistent decisions. A limiting factor for a broader adoption of AI technology is the inherent risks that come with giving up human control and oversight to “intelligent” machines. For sensitive tasks involving critical infrastructures and affecting human well-being or health, it is crucial to limit the possibility of improper, non-robust and unsafe decisions and actions. Before deploying an AI system, we see a strong need to validate its behavior, and thus establish guarantees that it will continue to perform as expected when deployed in a real-world environment. In pursuit of that objective, ways for humans to verify the agreement between the AI decision structure and their own ground-truth knowledge have been explored. Explainable AI (XAI) has developed as a subfield of AI, focused on exposing complex AI models to humans in a systematic and interpretable manner. The 22 chapters included in this book provide a timely snapshot of algorithms, theory, and applications of interpretable and explainable AI and AI techniques that have been proposed recently reflecting the current discourse in this field and providing directions of future development. The book is organized in six parts: towards AI transparency; methods for interpreting AI systems; explaining the decisions of AI systems; evaluating interpretability and explanations; applications of explainable AI; and software for explainable AI.

[Deep Learning with Structured Data](#) - Mark Ryan 2020-12-29

Deep Learning with Structured Data teaches you powerful data analysis techniques for tabular data and relational databases. Summary Deep learning offers the potential to identify complex patterns and relationships hidden in data of all sorts. Deep Learning with Structured Data shows you how to apply powerful deep learning analysis techniques to the kind of structured, tabular data you'll find in the relational databases that real-world businesses depend on. Filled with practical, relevant applications, this book teaches you how deep learning can augment your existing machine learning and business intelligence systems. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the technology Here's a dirty secret: Half of the time in most data science projects is spent cleaning and preparing data. But there's a better way: Deep learning techniques optimized for tabular data and relational databases deliver insights and analysis without requiring intense feature engineering. Learn the skills to unlock deep learning performance with much less data filtering, validating, and scrubbing. About the book Deep Learning with Structured Data teaches you powerful data analysis techniques for tabular data and relational databases. Get started using a dataset based on the Toronto transit system. As you work through the book, you'll learn how easy it is to set up tabular data for deep learning, while solving crucial production concerns like deployment and performance monitoring. What's inside When and where to use deep learning The architecture of a Keras deep learning model Training, deploying, and maintaining models Measuring performance About the reader For readers with intermediate Python and machine learning skills. About the author Mark Ryan is a Data Science Manager at Intact Insurance. He holds a Master's degree in Computer Science from the University of Toronto. Table of Contents 1 Why deep learning with structured data? 2 Introduction to the example problem and Pandas dataframes 3 Preparing the data, part 1: Exploring and cleansing the data 4 Preparing the data, part 2: Transforming the data 5 Preparing and building the model 6 Training the model and running experiments 7

More experiments with the trained model 8 Deploying the model 9 Recommended next steps

Deep Learning Applications - M. Arif Wani 2020-02-29

This book presents a compilation of selected papers from the 17th IEEE International Conference on Machine Learning and Applications (IEEE ICMLA 2018), focusing on use of deep learning technology in application like game playing, medical applications, video analytics, regression/classification, object detection/recognition and robotic control in industrial environments. It highlights novel ways of using deep neural networks to solve real-world problems, and also offers insights into deep learning architectures and algorithms, making it an essential reference guide for academic researchers, professionals, software engineers in industry, and innovative product developers.

[Learn Keras for Deep Neural Networks](#) - Jojo Moolayil 2018-12-07

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.

Grokking Deep Learning - Andrew W. Trask 2019-01-23

Summary Grokking Deep Learning teaches you to build deep learning neural networks from scratch! In his engaging style, seasoned deep learning expert Andrew Trask shows you the science under the hood, so you grok for yourself every detail of training neural networks. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology Deep learning, a branch of artificial intelligence, teaches computers to learn by using neural networks, technology inspired by the human brain. Online text translation, self-driving cars, personalized product recommendations, and virtual voice assistants are just a few of the exciting modern advancements possible thanks to deep learning. About the Book Grokking Deep Learning teaches you to build deep learning neural networks from scratch! In his engaging style, seasoned deep learning expert Andrew Trask shows you the science under

the hood, so you grok for yourself every detail of training neural networks. Using only Python and its math-supporting library, NumPy, you'll train your own neural networks to see and understand images, translate text into different languages, and even write like Shakespeare! When you're done, you'll be fully prepared to move on to mastering deep learning frameworks. What's inside The science behind deep learning Building and training your own neural networks Privacy concepts, including federated learning Tips for continuing your pursuit of deep learning About the Reader For readers with high school-level math and intermediate programming skills. About the Author Andrew Trask is a PhD student at Oxford University and a research scientist at DeepMind. Previously, Andrew was a researcher and analytics product manager at Digital Reasoning, where he trained the world's largest artificial neural network and helped guide the analytics roadmap for the Synthesys cognitive computing platform. Table of Contents Introducing deep learning: why you should learn it Fundamental concepts: how do machines learn? Introduction to neural prediction: forward propagation Introduction to neural learning: gradient descent Learning multiple weights at a time: generalizing gradient descent Building your first deep neural network: introduction to backpropagation How to picture neural networks: in your head and on paper Learning signal and ignoring noise: introduction to regularization and batching Modeling probabilities and nonlinearities: activation functions Neural learning about edges and corners: intro to convolutional neural networks Neural networks that understand language: king - man + woman == ? Neural networks that write like Shakespeare: recurrent layers for variable-length data Introducing automatic optimization: let's build a deep learning framework Learning to write like Shakespeare: long short-term memory Deep learning on unseen data: introducing federated learning Where to go from here: a brief guide

Efficient Processing of Deep Neural Networks - Vivienne Sze 2020-06-24

This book provides a structured treatment of the key principles and techniques for enabling efficient processing of deep neural networks (DNNs). DNNs are currently widely used for many artificial intelligence (AI) applications, including computer vision, speech recognition, and robotics. While DNNs deliver state-of-the-art accuracy on many AI tasks, it comes at the cost of high computational complexity. Therefore, techniques that enable efficient processing of deep neural networks to improve metrics—such as energy efficiency, throughput, and latency—without sacrificing accuracy or increasing hardware costs are critical to enabling the wide deployment of DNNs in AI systems. The book includes background on DNN processing; a description and taxonomy of hardware architectural approaches for designing DNN accelerators; key metrics for evaluating and comparing different designs; features of the DNN processing that are amenable to hardware/algorithm co-design to improve energy efficiency and throughput; and opportunities for applying new technologies. Readers will find a structured introduction to the field as well as a formalization and organization of key concepts from contemporary works that provides insights that may spark new ideas.

[MATLAB Deep Learning](#) - Phil Kim 2017-06-15

Get started with MATLAB for deep learning and AI with this in-depth primer. In this book, you start with machine learning fundamentals, then move on to neural networks, deep learning, and then convolutional neural networks. In a blend of fundamentals and applications, MATLAB Deep Learning employs MATLAB as the underlying programming language and tool for the examples and case studies in this book. With this book, you'll be able to tackle some of today's real world big data, smart bots, and other complex data problems. You'll see how deep learning is a complex and more intelligent aspect of machine learning for modern smart data analysis and usage. What You'll Learn Use MATLAB for deep learning Discover neural networks and multi-layer neural networks Work with convolution and pooling layers Build a MNIST example with these layers Who This Book Is For Those who want to learn deep learning using MATLAB. Some MATLAB experience may be useful.

Neural Networks and Deep Learning - Pat Nakamoto 2018-01-12

Ready to crank up a neural network to get your self-driving car pick up the kids from school? Want to add 'Deep Learning' to your LinkedIn profile? Well, hold on there... Before you embark on your epic journey into the world of deep learning, there is basic theory to march through first! Take a step-by-step journey through the basics of Neural Networks and Deep Learning, made so simple that...even your granny could understand it! What you will gain from this book: * A deep understanding of how a Neural Network and

Deep Learning work * A basics comprehension on how to build a Deep Neural Network from scratch Who this book is for: * Beginners who want to approach the topic, but are too afraid of complex math to start! What's Inside? * A brief introduction to Machine Learning * Two main Types of Machine Learning Algorithms * A practical example of Unsupervised Learning * What are Neural Networks? * McCulloch-Pitts's Neuron * Types of activation function * Types of network architectures * Learning processes * Advantages and disadvantages * Let us give a memory to our Neural Network * The example of book writing Software * Deep learning: the ability of learning to learn * How does Deep Learning work? * Main architectures and algorithms * Main types of DNN * Available Frameworks and libraries * Convolutional Neural Networks * Tunnel Vision * Convolution * The right Architecture for a Neural Network * Test your Neural Network * A general overview of Deep Learning * What are the limits of Deep Learning? * Deep Learning: the basics * Layers, Learning paradigms, Training, Validation * Main architectures and algorithms * Models for Deep Learning * Probabilistic graphic models * Restricted Boltzmann Machines * Deep Belief Networks * Available Frameworks and libraries * TensorFlow Hit download. Now!

Neural Networks and Deep Learning - Charu C. Aggarwal 2018-08-25

This book covers both classical and modern models in deep learning. The primary focus is on the theory and algorithms of deep learning. The theory and algorithms of neural networks are particularly important for understanding important concepts, so that one can understand the important design concepts of neural architectures in different applications. Why do neural networks work? When do they work better than off-the-shelf machine-learning models? When is depth useful? Why is training neural networks so hard? What are the pitfalls? The book is also rich in discussing different applications in order to give the practitioner a flavor of how neural architectures are designed for different types of problems. Applications associated with many different areas like recommender systems, machine translation, image captioning, image classification, reinforcement-learning based gaming, and text analytics are covered. The chapters of this book span three categories: The basics of neural networks: Many traditional machine learning models can be understood as special cases of neural networks. An emphasis is placed in the first two chapters on understanding the relationship between traditional machine learning and neural networks. Support vector machines, linear/logistic regression, singular value decomposition, matrix factorization, and recommender systems are shown to be special cases of neural networks. These methods are studied together with recent feature engineering methods like word2vec. Fundamentals of neural networks: A detailed discussion of training and regularization is provided in Chapters 3 and 4. Chapters 5 and 6 present radial-basis function (RBF) networks and restricted Boltzmann machines. Advanced topics in neural networks: Chapters 7 and 8 discuss recurrent neural networks and convolutional neural networks. Several advanced topics like deep reinforcement learning, neural Turing machines, Kohonen self-organizing maps, and generative adversarial networks are introduced in Chapters 9 and 10. The book is written for graduate students, researchers, and practitioners. Numerous exercises are available along with a solution manual to aid in classroom teaching. Where possible, an application-centric view is highlighted in order to provide an understanding of the practical uses of each class of techniques.

Deep Learning with Python - Francois Chollet 2017-11-30

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

[Learning Deep Architectures for AI](#) - Yoshua Bengio 2009

Theoretical results suggest that in order to learn the kind of complicated functions that can represent high-level abstractions (e.g. in vision, language, and other AI-level tasks), one may need deep architectures. Deep architectures are composed of multiple levels of non-linear operations, such as in neural nets with many hidden layers or in complicated propositional formulae re-using many sub-formulae. Searching the parameter space of deep architectures is a difficult task, but learning algorithms such as those for Deep Belief Networks have recently been proposed to tackle this problem with notable success, beating the state-of-the-art in certain areas. This paper discusses the motivations and principles regarding learning algorithms for deep architectures, in particular those exploiting as building blocks unsupervised learning of single-layer models such as Restricted Boltzmann Machines, used to construct deeper models such as Deep Belief Networks.

[Programming Machine Learning](#) - Paolo Perrotta 2020-03-31

You've decided to tackle machine learning - because you're job hunting, embarking on a new project, or just think self-driving cars are cool. But where to start? It's easy to be intimidated, even as a software developer. The good news is that it doesn't have to be that hard. Master machine learning by writing code one line at a time, from simple learning programs all the way to a true deep learning system. Tackle the hard topics by breaking them down so they're easier to understand, and build your confidence by getting your hands dirty. Peel away the obscurities of machine learning, starting from scratch and going all the way to deep learning. Machine learning can be intimidating, with its reliance on math and algorithms that most programmers don't encounter in their regular work. Take a hands-on approach, writing the Python code yourself, without any libraries to obscure what's really going on. Iterate on your design, and add layers of complexity as you go. Build an image recognition application from scratch with supervised learning. Predict the future with linear regression. Dive into gradient descent, a fundamental algorithm that drives most of machine learning. Create perceptrons to classify data. Build neural networks to tackle more complex and sophisticated data sets. Train and refine those networks with backpropagation and batching. Layer the neural networks, eliminate overfitting, and add convolution to transform your neural network into a true deep learning system. Start from the beginning and code your way to machine learning mastery. What You Need: The examples in this book are written in Python, but don't worry if you don't know this language: you'll pick up all the Python you need very quickly. Apart from that, you'll only need your computer, and your code-adept brain.

TensorFlow Deep Learning Projects - Alexey Grigorev 2018-03-28

Leverage the power of Tensorflow to design deep learning systems for a variety of real-world scenarios Key Features Build efficient deep learning pipelines using the popular Tensorflow framework Train neural networks such as ConvNets, generative models, and LSTMs Includes projects related to Computer Vision, stock prediction, chatbots and more Book Description TensorFlow is one of the most popular frameworks

used for machine learning and, more recently, deep learning. It provides a fast and efficient framework for training different kinds of deep learning models, with very high accuracy. This book is your guide to master deep learning with TensorFlow with the help of 10 real-world projects. TensorFlow Deep Learning Projects starts with setting up the right TensorFlow environment for deep learning. Learn to train different types of deep learning models using TensorFlow, including Convolutional Neural Networks, Recurrent Neural Networks, LSTMs, and Generative Adversarial Networks. While doing so, you will build end-to-end deep learning solutions to tackle different real-world problems in image processing, recommendation systems, stock prediction, and building chatbots, to name a few. You will also develop systems that perform machine translation, and use reinforcement learning techniques to play games. By the end of this book, you will have mastered all the concepts of deep learning and their implementation with TensorFlow, and will be able to build and train your own deep learning models with TensorFlow confidently. What you will learn Set up the TensorFlow environment for deep learning Construct your own ConvNets for effective image processing Use LSTMs for image caption generation Forecast stock prediction accurately with an LSTM architecture Learn what semantic matching is by detecting duplicate Quora questions Set up an AWS instance with TensorFlow to train GANs Train and set up a chatbot to understand and interpret human input Build an AI capable of playing a video game by itself -and win it! Who this book is for This book is for data scientists, machine learning developers as well as deep learning practitioners, who want to build interesting deep learning projects that leverage the power of Tensorflow. Some understanding of machine learning and deep learning, and familiarity with the TensorFlow framework is all you need to get started with this book.

Deep Learning for Biomedical Data Analysis - Mourad Elloumi 2021-07-13

This book is the first overview on Deep Learning (DL) for biomedical 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 overcrowding the reader's head. It presents the results of the latest investigations 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.

[Deep Learning](#) - John D. Kelleher 2019-09-10

An accessible introduction to the artificial intelligence technology that enables computer vision, speech recognition, machine translation, and driverless cars. Deep learning is an artificial intelligence technology that enables computer vision, speech recognition in mobile phones, machine translation, AI games, driverless cars, and other applications. When we use consumer products from Google, Microsoft, Facebook, Apple, or Baidu, we are often interacting with a deep learning system. In this volume in the MIT Press Essential Knowledge series, computer scientist John Kelleher offers an accessible and concise but comprehensive introduction to the fundamental technology at the heart of the artificial intelligence revolution. Kelleher explains that deep learning enables data-driven decisions by identifying and extracting patterns from large datasets; its ability to learn from complex data makes deep learning ideally suited to

take advantage of the rapid growth in big data and computational power. Kelleher also explains some of the basic concepts in deep learning, presents a history of advances in the field, and discusses the current state of the art. He describes the most important deep learning architectures, including autoencoders, recurrent neural networks, and long short-term networks, as well as such recent developments as Generative Adversarial Networks and capsule networks. He also provides a comprehensive (and comprehensible) introduction to the two fundamental algorithms in deep learning: gradient descent and backpropagation. Finally, Kelleher considers the future of deep learning—major trends, possible developments, and significant challenges.

Deep Learning from Scratch - Seth Weidman 2019-09-09

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 details of important 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

The Deep Learning Revolution - Terrence J. Sejnowski 2018-10-23

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 experience 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.

Deep Learning Illustrated - Jon Krohn 2019-08-05

"The authors' clear visual style provides a comprehensive look at what's currently possible with artificial neural networks as well as a glimpse of the magic that's to come." -Tim Urban, author of Wait But Why Fully Practical, Insightful Guide to Modern Deep Learning Deep learning is transforming software, facilitating powerful new artificial intelligence capabilities, and driving unprecedented algorithm performance. Deep Learning Illustrated is uniquely intuitive and offers a complete introduction to the discipline's techniques. Packed with full-color figures and easy-to-follow code, it sweeps away the complexity of building deep learning models, making the subject approachable and fun to learn. World-class instructor and practitioner Jon Krohn—with visionary content from Grant Beyleveld and beautiful illustrations by Aglaé Bassens—presents straightforward analogies to explain what deep learning is, why it has become so popular, and how it relates to other machine learning approaches. Krohn has created a practical reference and tutorial for developers, data scientists, researchers, analysts, and students who

want to start applying it. He illuminates theory with hands-on Python code in accompanying Jupyter notebooks. To help you progress quickly, he focuses on the versatile deep learning library Keras to nimbly construct efficient TensorFlow models; PyTorch, the leading alternative library, is also covered. You'll gain a pragmatic understanding of all major deep learning approaches and their uses in applications ranging from machine vision and natural language processing to image generation and game-playing algorithms. Discover what makes deep learning systems unique, and the implications for practitioners Explore new tools that make deep learning models easier to build, use, and improve Master essential theory: artificial neurons, training, optimization, convolutional nets, recurrent nets, generative adversarial networks (GANs), deep reinforcement learning, and more Walk through building interactive deep learning applications, and move forward with your own artificial intelligence projects Register your book for convenient access to downloads, updates, and/or corrections as they become available. See inside book for details.

Machine Learning For Dummies - John Paul Mueller 2021-02-09

One of Mark Cuban's top reads for better understanding A.I. (inc.com, 2021) Your comprehensive entry-level guide to machine learning While machine learning expertise doesn't quite mean you can create your own Turing Test-proof android—as in the movie Ex Machina—it is a form of artificial intelligence and one of the most exciting technological means of identifying opportunities and solving problems fast and on a large scale. Anyone who masters the principles of machine learning is mastering a big part of our tech future and opening up incredible new directions in careers that include fraud detection, optimizing search results, serving real-time ads, credit-scoring, building accurate and sophisticated pricing models—and way, way more. Unlike most machine learning books, the fully updated 2nd Edition of Machine Learning For Dummies doesn't assume you have years of experience using programming languages such as Python (R source is also included in a downloadable form with comments and explanations), but lets you in on the ground floor, covering the entry-level materials that will get you up and running building models you need to perform practical tasks. It takes a look at the underlying—and fascinating—math principles that power machine learning but also shows that you don't need to be a math whiz to build fun new tools and apply them to your work and study. Understand the history of AI and machine learning Work with Python 3.8 and TensorFlow 2.x (and R as a download) Build and test your own models Use the latest datasets, rather than the worn out data found in other books Apply machine learning to real problems Whether you want to learn for college or to enhance your business or career performance, this friendly beginner's guide is your best introduction to machine learning, allowing you to become quickly confident using this amazing and fast-developing technology that's impacting lives for the better all over the world.

Ultimate Step by Step Guide to Deep Learning Using Python - Daneyal Anis 2020-07-19

*Start your Data Science career using Python today!*Are you ready to start your new exciting career? Ready to master artificial intelligence and deep learning concepts?Are you overwhelmed with complexity of the books on this subject?Then let this breezy and fun little book on Python, Machine Learning and Deep Learning models make you a Data Scientist in 7 days!This book continues from where the first book in the series, Ultimate Step by Step Guide to Machine Learning Using Python, left of. In the first book you were introduced to Python concepts such as: -Data Structures like Pandas -Foundational libraries like Numpy, Seaborn and Scikit-Learn-Regression analysis-Classification-Clustering-Association Learning-Dimension ReductionThis book builds on those concepts to expand on Machine Learning algorithms like: -Linear and Logistical regression-Decision tree-Support vector machines (SVM)After that, this book takes you on a journey into Deep Learning and Neural Networks with important concepts and libraries like: -Convolutional and Recurrent Neural Networks-TensorFlow-Keras-PyTorch-Keras-Apache MXNet-Microsoft Cognitive Toolkit (CNTK)The final part of the book covers all foundational concepts that are required for Amazon Web Services (AWS) Certified Machine Learning Specialization by explaining how to deploy your models at scale on Cloud technologies. While AWS is used in the book for illustrative purposes, Microsoft Azure and Google Cloud are also introduced as alternative cloud technologies. After reading this book you will be able to: - Code in Python with confidence-Build new machine learning and deep learning models from scratch-Know how to clean and prepare your data for analytics-Speak confidently about statistical analysis techniquesData Science was ranked the fast-growing field by LinkedIn and Data Scientist is one of the most highly sought after and lucrative careers in the world!If you are on the fence about making the leap to a

new and lucrative career, this is the book for you! What sets this book apart from other books on the topic of Python and Machine learning: -Step by step code examples and explanation-Complex concepts explained visually-Real world applicability of the machine learning and deep learning models introduced-What do I need to get started? You will have a step by step action plan in place once you finish this book and finally feel that you, can master data science and artificial intelligence and start a lucrative and rewarding career! Ready to dive in to the exciting world of Python and Deep Learning? Then scroll up to the top and hit that BUY BUTTON!

[Deep Learning for Coders with fastai and PyTorch](#) - Jeremy Howard 2020-06-29

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.

Deep Learning - Andrew Glassner 2021-06-22

A richly-illustrated, full-color introduction to deep learning that offers visual and conceptual explanations instead of equations. You'll learn how to use key deep learning algorithms without the need for complex math. Ever since computers began beating us at chess, they've been getting better at a wide range of human activities, from writing songs and generating news articles to helping doctors provide healthcare. Deep learning is the source of many of these breakthroughs, and its remarkable ability to find patterns hiding in data has made it the fastest growing field in artificial intelligence (AI). Digital assistants on our phones use deep learning to understand and respond intelligently to voice commands; automotive systems use it to safely navigate road hazards; online platforms use it to deliver personalized suggestions for movies and books - the possibilities are endless. *Deep Learning: A Visual Approach* is for anyone who wants to understand this fascinating field in depth, but without any of the advanced math and programming usually required to grasp its internals. If you want to know how these tools work, and use them yourself, the answers are all within these pages. And, if you're ready to write your own programs, there are also plenty of supplemental Python notebooks in the accompanying Github repository to get you going. The book's conversational style, extensive color illustrations, illuminating analogies, and real-world examples expertly explain the key concepts in deep learning, including:

- How text generators create novel stories and articles
- How deep learning systems learn to play and win at human games
- How image classification systems identify objects or people in a photo
- How to think about probabilities in a way that's useful to everyday life
- How to use the machine learning techniques that form the core of modern AI

Intellectual adventurers of all kinds can use the powerful ideas covered in *Deep Learning: A Visual Approach* to build intelligent systems that help us better understand the world and everyone who lives in it. It's the future of AI, and this book allows you to fully envision it. Full Color Illustrations

TensorFlow for Deep Learning - Bharath Ramsundar 2018-03-01

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, understanding text, 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 professionals familiar with 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 hyperparameter optimization. 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.

Deep Learning with R for Beginners - Mark Hodnett 2019-05-17

Explore the world of neural networks by building powerful deep learning models using the R ecosystem. Key Features: Get to grips with the fundamentals of deep learning and neural networks. Use R 3.5 and its libraries and APIs to build deep learning models for computer vision and text processing. Implement effective deep learning systems in R with the help of end-to-end projects. Book Description: Deep learning finds practical applications in several domains, while R is the preferred language for designing and deploying deep learning models. This Learning Path introduces you to the basics of deep learning and even teaches you to build a neural network model from scratch. As you make your way through the chapters, you'll explore deep learning libraries and understand how to create deep learning models for a variety of challenges, right from anomaly detection to recommendation systems. The book will then help you cover advanced topics, such as generative adversarial networks (GANs), transfer learning, and large-scale deep learning in the cloud, in addition to model optimization, overfitting, and data augmentation. Through real-world projects, you'll also get up to speed with training convolutional neural networks (CNNs), recurrent neural networks (RNNs), and long short-term memory networks (LSTMs) in R. By the end of this Learning Path, you'll be well versed with deep learning and have the skills you need to implement a number of deep learning concepts in your research work or projects. This Learning Path includes content from the following Packt products: *R Deep Learning Essentials - Second Edition* by Joshua F. Wiley and Mark Hodnett. *R Deep Learning Projects* by Yuxi (Hayden) Liu and Pablo Maldonado. What you will learn: Implement credit card fraud detection with autoencoders. Train neural networks to perform handwritten digit recognition using MXNet. Reconstruct images using variational autoencoders. Explore the applications of autoencoder neural networks in clustering and dimensionality reduction. Create natural language processing (NLP) models using Keras and TensorFlow in R. Prevent models from overfitting the data to improve generalizability. Build shallow neural network prediction models. Who this book is for: This Learning Path is for aspiring data scientists, data analysts, machine learning developers, and deep learning enthusiasts who are well versed in machine learning concepts and are looking to explore the deep learning paradigm using R. A fundamental understanding of R programming and familiarity with the basic concepts of deep learning are necessary to get the most out of this Learning Path.

Applied Deep Learning - Umberto Michelucci 2018-09-07

Work with advanced topics in deep learning, such as optimization algorithms, hyper-parameter tuning, dropout, and error analysis as well as strategies to address typical problems encountered when training deep neural networks. You'll begin by studying the activation functions mostly with a single neuron (ReLU, sigmoid, and Swish), seeing how to perform linear and logistic regression using TensorFlow, and choosing the right cost function. The next section talks about more complicated neural network architectures with several layers and neurons and explores the problem of random initialization of weights. An entire chapter is dedicated to a complete overview of neural network error analysis, giving examples of solving problems originating from variance, bias, overfitting, and datasets coming from different distributions. *Applied Deep Learning* also discusses how to implement logistic regression completely from scratch without using any Python library except NumPy, to let you appreciate how libraries such as TensorFlow allow quick and efficient experiments. Case studies for each method are included to put into practice all theoretical information. You'll discover tips and tricks for writing optimized Python code (for example vectorizing loops with NumPy). What You Will Learn: Implement advanced techniques in the right way in Python and TensorFlow. Debug and optimize advanced methods (such as dropout and regularization). Carry out error analysis (to realize if one has a bias problem, a variance problem, a data offset problem, and so on). Set up a machine learning project focused on deep learning on a complex dataset. Who This Book Is For: Readers with a medium understanding of machine learning, linear algebra, calculus, and basic Python programming.

[Deep Learning](#) - Ian Goodfellow 2016-11-10

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.

Interpretable Machine Learning - Christoph Molnar 2020

This book is about making machine learning models and their decisions interpretable. After exploring the concepts of interpretability, you will learn about simple, interpretable models such as decision trees, decision rules and linear regression. Later chapters focus on general model-agnostic methods for interpreting black box models like feature importance and accumulated local effects and explaining individual predictions with Shapley values and LIME. All interpretation methods are explained in depth and discussed critically. How do they work under the hood? What are their strengths and weaknesses? How can their outputs be interpreted? This book will enable you to select and correctly apply the interpretation method that is most suitable for your machine learning project.

Deep learning: deep learning explained to your granny - a guide for beginners - PAT NAKAMOTO 2018-07-17

Ready to crank up a deep neural network to get your self-driving car pick up the kids from school? Want to add 'Deep Learning' to your LinkedIn profile? Well, hold on there... Before you embark on your epic journey into the world of deep learning, there is basic theory to march through first! Take a step-by-step journey through the basics of Deep Learning, made so simple that...even your granny could understand it! What you will gain from this book: * A deep understanding of how Deep Learning works * A basics comprehension on how to build a Deep Neural Network from scratch Who this book is for: * Beginners who want to approach the topic, but are too afraid of complex math to start! What's Inside? * A general overview of Deep Learning * What are the limits of Deep Learning? * Deep Learning: the basics * Layers, Learning paradigms, Training, Validation * Main architectures and algorithms * Convolutional Neural Networks * Models for Deep Learning * Probabilistic graphic models * Restricted Boltzmann Machines * Deep Belief Networks * Available Frameworks and libraries * TensorFlow Hit download. Now!

Deep Learning For Dummies - John Paul Mueller 2019-04-15

Take a deep dive into deep learning Deep learning provides the means for discerning patterns in the data that drive online business and social media outlets. Deep Learning for Dummies gives you the information you need to take the mystery out of the topic—and all of the underlying technologies associated with it. In no time, you'll make sense of those increasingly confusing algorithms, and find a simple and safe environment to experiment with deep learning. The book develops a sense of precisely what deep learning can do at a high level and then provides examples of the major deep learning application types. Includes sample code Provides real-world examples within the approachable text Offers hands-on activities to make

learning easier Shows you how to use Deep Learning more effectively with the right tools This book is perfect for those who want to better understand the basis of the underlying technologies that we use each and every day.

Deep Learning with R - Abhijit Ghatak 2019-04-13

Deep Learning with R introduces deep learning and neural networks using the R programming language. The book builds on the understanding of the theoretical and mathematical constructs and enables the reader to create applications on computer vision, natural language processing and transfer learning. The book starts with an introduction to machine learning and moves on to describe the basic architecture, different activation functions, forward propagation, cross-entropy loss and backward propagation of a simple neural network. It goes on to create different code segments to construct deep neural networks. It discusses in detail the initialization of network parameters, optimization techniques, and some of the common issues surrounding neural networks such as dealing with NaNs and the vanishing/exploding gradient problem. Advanced variants of multilayered perceptrons namely, convolutional neural networks and sequence models are explained, followed by application to different use cases. The book makes extensive use of the Keras and TensorFlow frameworks.

Handbook of Research on Deep Learning-Based Image Analysis Under Constrained and Unconstrained Environments - Raj, Alex Noel Joseph 2020-12-25

Recent advancements in imaging techniques and image analysis has broadened the horizons for their applications in various domains. Image analysis has become an influential technique in medical image analysis, optical character recognition, geology, remote sensing, and more. However, analysis of images under constrained and unconstrained environments require efficient representation of the data and complex models for accurate interpretation and classification of data. Deep learning methods, with their hierarchical/multilayered architecture, allow the systems to learn complex mathematical models to provide improved performance in the required task. The Handbook of Research on Deep Learning-Based Image Analysis Under Constrained and Unconstrained Environments provides a critical examination of the latest advancements, developments, methods, systems, futuristic approaches, and algorithms for image analysis and addresses its challenges. Highlighting concepts, methods, and tools including convolutional neural networks, edge enhancement, image segmentation, machine learning, and image processing, the book is an essential and comprehensive reference work for engineers, academicians, researchers, and students.

Artificial Intelligence with Python - Prateek Joshi 2017-01-27

Build real-world Artificial Intelligence applications with Python to intelligently interact with the world around you About This Book Step into the amazing world of intelligent apps using this comprehensive guide Enter the world of Artificial Intelligence, explore it, and create your own applications Work through simple yet insightful examples that will get you up and running with Artificial Intelligence in no time Who This Book Is For This book is for Python developers who want to build real-world Artificial Intelligence applications. This book is friendly to Python beginners, but being familiar with Python would be useful to play around with the code. It will also be useful for experienced Python programmers who are looking to use Artificial Intelligence techniques in their existing technology stacks. What You Will Learn Realize different classification and regression techniques Understand the concept of clustering and how to use it to automatically segment data See how to build an intelligent recommender system Understand logic programming and how to use it Build automatic speech recognition systems Understand the basics of heuristic search and genetic programming Develop games using Artificial Intelligence Learn how reinforcement learning works Discover how to build intelligent applications centered on images, text, and time series data See how to use deep learning algorithms and build applications based on it In Detail Artificial Intelligence is becoming increasingly relevant in the modern world where everything is driven by technology and data. It is used extensively across many fields such as search engines, image recognition, robotics, finance, and so on. We will explore various real-world scenarios in this book and you'll learn about various algorithms that can be used to build Artificial Intelligence applications. During the course of this book, you will find out how to make informed decisions about what algorithms to use in a given context. Starting from the basics of Artificial Intelligence, you will learn how to develop various building blocks using different data mining techniques. You will see how to implement different algorithms to get the best

possible results, and will understand how to apply them to real-world scenarios. If you want to add an intelligence layer to any application that's based on images, text, stock market, or some other form of data, this exciting book on Artificial Intelligence will definitely be your guide! Style and approach This highly practical book will show you how to implement Artificial Intelligence. The book provides multiple examples enabling you to create smart applications to meet the needs of your organization. In every chapter, we explain an algorithm, implement it, and then build a smart application.

Python Deep Learning - Valentino Zocca 2017-04-28

Take your machine learning skills to the next level by mastering Deep Learning concepts and algorithms using Python. About This Book Explore and create intelligent systems using cutting-edge deep learning techniques Implement deep learning algorithms and work with revolutionary libraries in Python Get real-world examples and easy-to-follow tutorials on Theano, TensorFlow, H2O and more Who This Book Is For This book is for Data Science practitioners as well as aspirants who have a basic foundational understanding of Machine Learning concepts and some programming experience with Python. A mathematical background with a conceptual understanding of calculus and statistics is also desired. What You Will Learn Get a practical deep dive into deep learning algorithms Explore deep learning further with Theano, Caffe, Keras, and TensorFlow Learn about two of the most powerful techniques at the core of many practical deep learning implementations: Auto-Encoders and Restricted Boltzmann Machines Dive into Deep Belief Nets and Deep Neural Networks Discover more deep learning algorithms with Dropout and Convolutional Neural Networks Get to know device strategies so you can use deep learning algorithms and libraries in the real world In Detail With an increasing interest in AI around the world, deep learning has attracted a great deal of public attention. Every day, deep learning algorithms are used broadly across different industries. The book will give you all the practical information available on the subject, including the best practices, using real-world use cases. You will learn to recognize and extract information to increase predictive accuracy and optimize results. Starting with a quick recap of important machine learning concepts, the book will delve straight into deep learning principles using Sci-kit learn. Moving ahead, you will learn to use the latest open source libraries such as Theano, Keras, Google's TensorFlow, and H2O. Use this guide to uncover the difficulties of pattern recognition, scaling data with greater accuracy and discussing deep learning algorithms and techniques. Whether you want to dive deeper into Deep Learning, or want to investigate how to get more out of this powerful technology, you'll find everything inside. Style and approach Python Machine Learning by example follows practical hands on approach. It walks you through the key elements of Python and its powerful machine learning libraries with the help of real world projects.

Deep Learning Essentials - Anurag Bhardwaj 2018-01-30

Get to grips with the essentials of deep learning by leveraging the power of Python Key Features Your one-stop solution to get started with the essentials of deep learning and neural network modeling Train different kinds of neural networks to tackle various problems in Natural Language Processing, computer vision, speech recognition, and more Covers popular Python libraries such as Tensorflow, Keras, and more, along with tips on training, deploying and optimizing your deep learning models in the best possible manner Book Description Deep Learning a trending topic in the field of Artificial Intelligence today and can be considered to be an advanced form of machine learning, which is quite tricky to master. This book will help you take your first steps in training efficient deep learning models and applying them in various practical scenarios. You will model, train, and deploy different kinds of neural networks such as Convolutional Neural Network, Recurrent Neural Network, and will see some of their applications in real-world domains including computer vision, natural language processing, speech recognition, and so on. You will build practical projects such as chatbots, implement reinforcement learning to build smart games, and develop expert systems for image captioning and processing. Popular Python library such as TensorFlow is used in this book to build the models. This book also covers solutions for different problems you might come across while training models, such as noisy datasets, small datasets, and more. This book does not assume any prior knowledge of deep learning. By the end of this book, you will have a firm understanding of the basics of deep learning and neural network modeling, along with their practical applications. What you will learn Get to grips with the core concepts of deep learning and neural networks Set up deep learning library

such as TensorFlow Fine-tune your deep learning models for NLP and Computer Vision applications Unify different information sources, such as images, text, and speech through deep learning Optimize and fine-tune your deep learning models for better performance Train a deep reinforcement learning model that plays a game better than humans Learn how to make your models get the best out of your GPU or CPU Who this book is for Aspiring data scientists and machine learning experts who have limited or no exposure to deep learning will find this book to be very useful. If you are looking for a resource that gets you up and running with the fundamentals of deep learning and neural networks, this book is for you. As the models in the book are trained using the popular Python-based libraries such as Tensorflow and Keras, it would be useful to have sound programming knowledge of Python.

Dive Into Deep Learning - Joanne Quinn 2019-07-15

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.

R Deep Learning Projects - Yuxi (Hayden) Liu 2018-02-22

5 real-world projects to help you master deep learning concepts Key Features Master the different deep learning paradigms and build real-world projects related to text generation, sentiment analysis, fraud detection, and more Get to grips with R's impressive range of Deep Learning libraries and frameworks such as deepnet, MXNetR, Tensorflow, H2O, Keras, and text2vec Practical projects that show you how to implement different neural networks with helpful tips, tricks, and best practices Book Description R is a popular programming language used by statisticians and mathematicians for statistical analysis, and is popularly used for deep learning. Deep Learning, as we all know, is one of the trending topics today, and is finding practical applications in a lot of domains. This book demonstrates end-to-end implementations of five real-world projects on popular topics in deep learning such as handwritten digit recognition, traffic light detection, fraud detection, text generation, and sentiment analysis. You'll learn how to train effective neural networks in R—including convolutional neural networks, recurrent neural networks, and LSTMs—and apply them in practical scenarios. The book also highlights how neural networks can be trained using GPU capabilities. You will use popular R libraries and packages—such as MXNetR, H2O, deepnet, and more—to implement the projects. By the end of this book, you will have a better understanding of deep learning concepts and techniques and how to use them in a practical setting. What you will learn Instrument Deep Learning models with packages such as deepnet, MXNetR, Tensorflow, H2O, Keras, and text2vec Apply neural networks to perform handwritten digit recognition using MXNet Get the knack of CNN models, Neural Network API, Keras, and TensorFlow for traffic sign classification - Implement credit card fraud detection with Autoencoders Master reconstructing images using variational autoencoders Wade through sentiment analysis from movie reviews Run from past to future and vice versa with bidirectional Long Short-Term Memory (LSTM) networks Understand the applications of Autoencoder Neural Networks in clustering and dimensionality reduction Who this book is for Machine learning professionals and data scientists looking to master deep learning by implementing practical projects in R

will find this book a useful resource. A knowledge of R programming and the basic concepts of deep learning is required to get the best out of this book.

Python Programming, Deep Learning - Anthony Adams 2021-12-17

Easily Boost Your Skills In Python Programming & Become A Master In Deep Learning & Data Analysis! Python is an interpreted, high-level, general-purpose programming language that emphasizes code readability with its notable use of significant whitespace. What makes Python so popular in the IT industry is that it uses an object-oriented approach, which enables programmers to write clear, logical code for all types of projects, whether big or small. Hone your Python Programming skills and gain a sharp edge over other programmers the EASIEST way possible... with this practical beginner's guide! In his 3-in-1 Python crash course for beginners, Anthony Adams gives novices like you simple, yet efficient tips and tricks to become a MASTER in Python coding for artificial intelligence, neural networks, machine learning, and data science/analysis! Here's what you'll get: Highly innovative ways to boost your understanding of Python programming, data analysis, and machine learning Quickly and effectively stop fraud with machine learning Practical and efficient exercises that make understanding Python quick & easy And so much more! As a beginner, you might feel a bit intimidated by the complexities of coding. Add the fact that most Python Programming crash course guides make learning harder than it has to be! With the help of this 3-in-1 guide, you will be given carefully sequenced Python Programming lessons that'll maximize your understanding, and equip you with all the skills for real-life application! Thrive in the IT industry with this comprehensive Python Programming crash course! Scroll up, Click on "Buy Now", and Start Learning Today!

Deep Learning - Josh Patterson 2017-07-28

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

Advanced Deep Learning with R - Bharatendra Rai 2019-12-17

Discover best practices for choosing, building, training, and improving deep learning models using Keras-R, and TensorFlow-R libraries Key Features Implement deep learning algorithms to build AI models with the help of tips and tricks Understand how deep learning models operate using expert techniques Apply reinforcement learning, computer vision, GANs, and NLP using a range of datasets Book Description Deep learning is a branch of machine learning based on a set of algorithms that attempt to model high-level

abstractions in data. Advanced Deep Learning with R will help you understand popular deep learning architectures and their variants in R, along with providing real-life examples for them. This deep learning book starts by covering the essential deep learning techniques and concepts for prediction and classification. You will learn about neural networks, deep learning architectures, and the fundamentals for implementing deep learning with R. The book will also take you through using important deep learning libraries such as Keras-R and TensorFlow-R to implement deep learning algorithms within applications. You will get up to speed with artificial neural networks, recurrent neural networks, convolutional neural networks, long short-term memory networks, and more using advanced examples. Later, you'll discover how to apply generative adversarial networks (GANs) to generate new images; autoencoder neural networks for image dimension reduction, image de-noising and image correction and transfer learning to prepare, define, train, and model a deep neural network. By the end of this book, you will be ready to implement your knowledge and newly acquired skills for applying deep learning algorithms in R through real-world examples. What you will learn Learn how to create binary and multi-class deep neural network models Implement GANs for generating new images Create autoencoder neural networks for image dimension reduction, image de-noising and image correction Implement deep neural networks for performing efficient text classification Learn to define a recurrent convolutional network model for classification in Keras Explore best practices and tips for performance optimization of various deep learning models Who this book is for This book is for data scientists, machine learning practitioners, deep learning researchers and AI enthusiasts who want to develop their skills and knowledge to implement deep learning techniques and algorithms using the power of R. A solid understanding of machine learning and working knowledge of the R programming language are required.

Deep Learning - Siddhartha Bhattacharyya 2020-06-22

This book focuses on the fundamentals of deep learning along with reporting on the current state-of-art research on deep learning. In addition, it provides an insight of deep neural networks in action with illustrative coding examples. Deep learning is a new area of machine learning research which has been introduced with the objective of moving ML closer to one of its original goals, i.e. artificial intelligence. Deep learning was developed as an ML approach to deal with complex input-output mappings. While traditional methods successfully solve problems where final value is a simple function of input data, deep learning techniques are able to capture composite relations between non-immediately related fields, for example between air pressure recordings and English words, millions of pixels and textual description, brand-related news and future stock prices and almost all real world problems. Deep learning is a class of nature inspired machine learning algorithms that uses a cascade of multiple layers of nonlinear processing units for feature extraction and transformation. Each successive layer uses the output from the previous layer as input. The learning may be supervised (e.g. classification) and/or unsupervised (e.g. pattern analysis) manners. These algorithms learn multiple levels of representations that correspond to different levels of abstraction by resorting to some form of gradient descent for training via backpropagation. Layers that have been used in deep learning include hidden layers of an artificial neural network and sets of propositional formulas. They may also include latent variables organized layer-wise in deep generative models such as the nodes in deep belief networks and deep boltzmann machines. Deep learning is part of state-of-the-art systems in various disciplines, particularly computer vision, automatic speech recognition (ASR) and human action recognition.