

Understanding Deep Convolutional Neural Networks With A

When people should go to the ebook stores, search establishment by shop, shelf by shelf, it is essentially problematic. This is why we offer the books compilations in this website. It will enormously ease you to see guide **understanding deep convolutional neural networks with a** as you such as.

By searching the title, publisher, or authors of guide you really want, you can discover them rapidly. In the house, workplace, or perhaps in your method can be all best place within net connections. If you point toward to download and install the understanding deep convolutional neural networks with a, it is utterly simple then, previously currently we extend the colleague to purchase and create bargains to download and install understanding deep convolutional neural networks with a fittingly simple!

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.

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.

Transparency and Interpretability for Learned Representations of Artificial Neural Networks -

Richard Meyes 2022-12-28

Artificial intelligence (AI) is a concept, whose meaning and perception has changed considerably over the last decades. Starting off with individual and purely theoretical research efforts in the 1950s, AI has grown into a fully developed research field of modern times and may arguably emerge as one of the most important technological advancements of mankind. Despite these rapid technological advancements, some key questions revolving around the matter of transparency, interpretability and explainability of an AI's decision-making remain unanswered. Thus, a young research field coined with the general term Explainable AI (XAI) has emerged from increasingly strict requirements for AI to be used in safety critical or ethically sensitive domains. An important research branch of XAI is to develop methods that help to facilitate a

deeper understanding for the learned knowledge of artificial neural systems. In this book, a series of scientific studies are presented that shed light on how to adopt an empirical neuroscience inspired approach to investigate a neural network's learned representation in the same spirit as neuroscientific studies of the brain.

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 - Frank Millstein 2018-03-21
Deep Learning - 2 BOOK BUNDLE!! Deep Learning with Keras This book will introduce you to various supervised and unsupervised deep learning algorithms like the multilayer perceptron, linear regression and other more advanced deep convolutional and recurrent neural networks. You will also learn about image processing, handwritten recognition, object recognition and much more. Furthermore, you will get familiar with recurrent neural networks like LSTM and GAN as you explore processing sequence data like time series, text, and audio. The book will definitely be your best companion on this great deep learning journey with Keras introducing you to the basics you need to know in order to take next steps and learn more advanced deep neural networks. Here Is a Preview of What You'll Learn Here... The difference between deep learning and machine learning Deep neural networks Convolutional neural networks Building deep learning models

with Keras Multi-layer perceptron network models Activation functions Handwritten recognition using MNIST Solving multi-class classification problems Recurrent neural networks and sequence classification And much more... Convolutional Neural Networks in Python This book covers the basics behind Convolutional Neural Networks by introducing you to this complex world of deep learning and artificial neural networks in a simple and easy to understand way. It is perfect for any beginner out there looking forward to learning more about this machine learning field. This book is all about how to use convolutional neural networks for various image, object and other common classification problems in Python. Here, we also take a deeper look into various Keras layer used for building CNNs we take a look at different activation functions and much more, which will eventually lead you to creating highly accurate models able of performing great task results on various image classification, object

classification and other problems. Therefore, at the end of the book, you will have a better insight into this world, thus you will be more than prepared to deal with more complex and challenging tasks on your own. Here Is a Preview of What You'll Learn In This Book... Convolutional neural networks structure How convolutional neural networks actually work Convolutional neural networks applications The importance of convolution operator Different convolutional neural networks layers and their importance Arrangement of spatial parameters How and when to use stride and zero-padding Method of parameter sharing Matrix multiplication and its importance Pooling and dense layers Introducing non-linearity relu activation function How to train your convolutional neural network models using backpropagation How and why to apply dropout CNN model training process How to build a convolutional neural network Generating predictions and calculating loss functions How

to train and evaluate your MNIST classifier How to build a simple image classification CNN And much, much more! Get this book bundle NOW and SAVE money!

Deep Learning: Concepts and Architectures

Witold Pedrycz 2019-10-29

This book introduces readers to the fundamental concepts of deep learning and offers practical insights into how this learning paradigm supports automatic mechanisms of structural knowledge representation. It discusses a number of multilayer architectures giving rise to tangible and functionally meaningful pieces of knowledge, and shows how the structural developments have become essential to the successful delivery of competitive practical solutions to real-world problems. The book also demonstrates how the architectural developments, which arise in the setting of deep learning, support detailed learning and refinements to the system design. Featuring detailed descriptions of the current trends in the

design and analysis of deep learning topologies, the book offers practical guidelines and presents competitive solutions to various areas of language modeling, graph representation, and forecasting.

Convolutional Neural Networks in Visual Computing - Ragav Venkatesan 2017-10-23

This book covers the fundamentals in designing and deploying techniques using deep architectures. It is intended to serve as a beginner's guide to engineers or students who want to have a quick start on learning and/or building deep learning systems. This book provides a good theoretical and practical understanding and a complete toolkit of basic information and knowledge required to understand and build convolutional neural networks (CNN) from scratch. The book focuses explicitly on convolutional neural networks, filtering out other material that co-occur in many deep learning books on CNN topics.

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.

Introduction to Deep Learning Business Applications for Developers - Armando Vieira
2018-05-02

Discover the potential applications, challenges, and opportunities of deep learning from a business perspective with technical examples. These applications include image recognition, segmentation and annotation, video processing and annotation, voice recognition, intelligent personal assistants, automated translation, and autonomous vehicles. An Introduction to Deep Learning Business Applications for Developers covers some common DL algorithms such as content-based recommendation algorithms and natural language processing. You'll explore examples, such as video prediction with fully convolutional neural networks (FCNN) and residual neural networks (ResNets). You will also see applications of DL for controlling

robotics, exploring the DeepQ learning algorithm with Monte Carlo Tree search (used to beat humans in the game of Go), and modeling for financial risk assessment. There will also be mention of the powerful set of algorithms called Generative Adversarial Neural networks (GANs) that can be applied for image colorization, image completion, and style transfer. After reading this book you will have an overview of the exciting field of deep neural networks and an understanding of most of the major applications of deep learning. The book contains some coding examples, tricks, and insights on how to train deep learning models using the Keras framework. What You Will Learn Find out about deep learning and why it is so powerful Work with the major algorithms available to train deep learning models See the major breakthroughs in terms of applications of deep learning Run simple examples with a selection of deep learning libraries Discover the areas of impact of deep learning in business Who This Book Is For

Data scientists, entrepreneurs, and business developers.

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 - Pat Nakamoto 2018-06-30

What's Inside? This includes 3 manuscripts:
Book 1: Neural Networks & Deep Learning:
Deep Learning explained to your granny - A visual introduction for beginners who want to make their own Deep Learning Neural Network... 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! * 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 Book2: Deep Learning: Deep Learning explained to your granny - A guide for Beginners... 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 Book 3: Big Data: The revolution that is transforming our work, market and world... "Within 2 days we produce the same amount of data generated by at the beginning of the civilization until 2003," said Eric Schmidt in 2010. According to IBM, by 2020 the world will have generated a mass of data on the order of 40 zettabyte (10²¹Byte). Just think, for example, of digital content such as photos, videos, blogs, posts, and everything that revolves around social networks; only Facebook marks 30 billion pieces of content each month shared by its users. The explosion of social networks, combined with the emergence of smartphones, justifies the fact that one of the recurring terms of recent years in the field of innovation, marketing and IT is "Big Data." The term Big Data indicates data produced in massive quantities, with remarkable rapidity and in the most diverse formats, which

require technologies and resources that go far beyond conventional data management and storage systems. In order to obtain from the use of this data the maximum results in the shortest possible time or even in real time, specific tools with high computing capabilities are necessary. But what does the Big Data phenomenon mean? Is the proliferation of data simply the sign of an increasingly invasive world? Or is there something more to it? Pat Nakamoto will guide you through the discovery of the world of Big data, which, according to experts, in the near future could become the new gold or oil, in what is a real Data Driven economy.

Deep Neural Networks and Data for Automated Driving - Tim Fingscheidt

2022-07-19

This open access book brings together the latest developments from industry and research on automated driving and artificial intelligence. Environment perception for highly automated driving heavily employs deep neural networks,

facing many challenges. How much data do we need for training and testing? How to use synthetic data to save labeling costs for training? How do we increase robustness and decrease memory usage? For inevitably poor conditions: How do we know that the network is uncertain about its decisions? Can we understand a bit more about what actually happens inside neural networks? This leads to a very practical problem particularly for DNNs employed in automated driving: What are useful validation techniques and how about safety? This book unites the views from both academia and industry, where computer vision and machine learning meet environment perception for highly automated driving. Naturally, aspects of data, robustness, uncertainty quantification, and, last but not least, safety are at the core of it. This book is unique: In its first part, an extended survey of all the relevant aspects is provided. The second part contains the detailed technical elaboration of the various questions mentioned above.

A Guide to Convolutional Neural Networks for Computer Vision Salman Khan 2018-02-13
Computer vision has become increasingly important and effective in recent years due to its wide-ranging applications in areas as diverse as smart surveillance and monitoring, health and medicine, sports and recreation, robotics, drones, and self-driving cars. Visual recognition tasks, such as image classification, localization, and detection, are the core building blocks of many of these applications, and recent developments in Convolutional Neural Networks (CNNs) have led to outstanding performance in these state-of-the-art visual recognition tasks and systems. As a result, CNNs now form the crux of deep learning algorithms in computer vision. This self-contained guide will benefit those who seek to both understand the theory behind CNNs and to gain hands-on experience on the application of CNNs in computer vision. It provides a comprehensive introduction to CNNs starting with the essential concepts behind

neural networks: training, regularization, and optimization of CNNs. The book also discusses a wide range of loss functions, network layers, and popular CNN architectures, reviews the different techniques for the evaluation of CNNs, and presents some popular CNN tools and libraries that are commonly used in computer vision. Further, this text describes and discusses case studies that are related to the application of CNN in computer vision, including image classification, object detection, semantic segmentation, scene understanding, and image generation. This book is ideal for undergraduate and graduate students, as no prior background knowledge in the field is required to follow the material, as well as new researchers, developers, engineers, and practitioners who are interested in gaining a quick understanding of CNN models.

Deep Learning with Python Nikhil Ketkar
2021-04-10

Master the practical aspects of implementing

deep learning solutions with PyTorch, using a hands-on approach to understanding both theory and practice. This updated edition will prepare you for applying deep learning to real world problems with a sound theoretical foundation and practical know-how with PyTorch, a platform developed by Facebook's Artificial Intelligence Research Group. You'll start with a perspective on how and why deep learning with PyTorch has emerged as an path-breaking framework with a set of tools and techniques to solve real-world problems. Next, the book will ground you with the mathematical fundamentals of linear algebra, vector calculus, probability and optimization. Having established this foundation, you'll move on to key components and functionality of PyTorch including layers, loss functions and optimization algorithms. You'll also gain an understanding of Graphical Processing Unit (GPU) based computation, which is essential for training deep learning models. All the key architectures in deep

learning are covered, including feedforward networks, convolution neural networks, recurrent neural networks, long short-term memory networks, autoencoders and generative adversarial networks. Backed by a number of tricks of the trade for training and optimizing deep learning models, this edition of Deep Learning with Python explains the best practices in taking these models to production with PyTorch. What You'll Learn Review machine learning fundamentals such as overfitting, underfitting, and regularization. Understand deep learning fundamentals such as feed-forward networks, convolution neural networks, recurrent neural networks, automatic differentiation, and stochastic gradient descent. Apply in-depth linear algebra with PyTorch Explore PyTorch fundamentals and its building blocks Work with tuning and optimizing models Who This Book Is For Beginners with a working knowledge of Python who want to understand Deep Learning in a practical, hands-on manner.

Deep Learning: Convergence to Big Data Analytics - Murad Khan 2018-12-30

This book presents deep learning techniques, concepts, and algorithms to classify and analyze big data. Further, it offers an introductory level understanding of the new programming languages and tools used to analyze big data in real-time, such as Hadoop, SPARK, and GRAPHX. Big data analytics using traditional techniques face various challenges, such as fast, accurate and efficient processing of big data in real-time. In addition, the Internet of Things is progressively increasing in various fields, like smart cities, smart homes, and e-health. As the enormous number of connected devices generate huge amounts of data every day, we need sophisticated algorithms to deal, organize, and classify this data in less processing time and space. Similarly, existing techniques and algorithms for deep learning in big data field have several advantages thanks to the two main branches of the deep learning, i.e. convolution

and deep belief networks. This book offers insights into these techniques and applications based on these two types of deep learning. Further, it helps students, researchers, and newcomers understand big data analytics based on deep learning approaches. It also discusses various machine learning techniques in concatenation with the deep learning paradigm to support high-end data processing, data classifications, and real-time data processing issues. The classification and presentation are kept quite simple to help the readers and students grasp the basics concepts of various deep learning paradigms and frameworks. It mainly focuses on theory rather than the mathematical background of the deep learning concepts. The book consists of 5 chapters, beginning with an introductory explanation of big data and deep learning techniques, followed by integration of big data and deep learning techniques and lastly the future directions. [Learning TensorFlow](#) - Tom Hope 2017-08-09

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

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

[Deep Learning in Mining of Visual Content](#) - Akka Zemhari 2020-01-22

This book provides the reader with the fundamental knowledge in the area of deep learning with application to visual content mining. The authors give a fresh view on Deep learning approaches both from the point of view of image understanding and supervised machine learning. It contains chapters which introduce theoretical and mathematical foundations of

neural networks and related optimization methods. Then it discusses some particular very popular architectures used in the domain: convolutional neural networks and recurrent neural networks. Deep Learning is currently at the heart of most cutting edge technologies. It is in the core of the recent advances in Artificial Intelligence. Visual information in Digital form is constantly growing in volume. In such active domains as Computer Vision and Robotics visual information understanding is based on the use of deep learning. Other chapters present applications of deep learning for visual content mining. These include attention mechanisms in deep neural networks and application to digital cultural content mining. An additional application field is also discussed, and illustrates how deep learning can be of very high interest to computer-aided diagnostics of Alzheimer's disease on multimodal imaging. This book targets advanced-level students studying computer science including computer vision,

data analytics and multimedia. Researchers and professionals working in computer science, signal and image processing may also be interested in this book.

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 Applied Deep Learning* Umberto Michelucci 2019-09-28

Develop and optimize deep learning models with advanced architectures. This book teaches you the intricate details and subtleties of the algorithms that are at the core of convolutional neural networks. In *Advanced Applied Deep Learning*, you will study advanced topics on CNN and object detection using Keras and

TensorFlow. Along the way, you will look at the fundamental operations in CNN, such as convolution and pooling, and then look at more advanced architectures such as inception networks, resnets, and many more. While the book discusses theoretical topics, you will discover how to work efficiently with Keras with many tricks and tips, including how to customize logging in Keras with custom callback classes, what is eager execution, and how to use it in your models. Finally, you will study how object detection works, and build a complete implementation of the YOLO (you only look once) algorithm in Keras and TensorFlow. By the end of the book you will have implemented various models in Keras and learned many advanced tricks that will bring your skills to the next level. What You Will Learn See how convolutional neural networks and object detection work Save weights and models on disk Pause training and restart it at a later stage Use hardware acceleration (GPUs) in your code Work with the

Dataset TensorFlow abstraction and use pre-trained models and transfer learning Remove and add layers to pre-trained networks to adapt them to your specific project Apply pre-trained models such as Alexnet and VGG16 to new datasets Who This Book Is For Scientists and researchers with intermediate-to-advanced Python and machine learning know-how. Additionally, intermediate knowledge of Keras and TensorFlow is expected. *Geometry of Deep Learning* Jong Chul Ye 2022 The focus of this book is on providing students with insights into geometry that can help them understand deep learning from a unified perspective. Rather than describing deep learning as an implementation technique, as is usually the case in many existing deep learning books, here, deep learning is explained as an ultimate form of signal processing techniques that can be imagined. To support this claim, an overview of classical kernel machine learning approaches is presented, and their advantages and limitations are explained. Following a

detailed explanation of the basic building blocks of deep neural networks from a biological and algorithmic point of view, the latest tools such as attention, normalization, Transformer, BERT, GPT-3, and others are described. Here, too, the focus is on the fact that in these heuristic approaches, there is an important, beautiful geometric structure behind the intuition that enables a systematic understanding. A unified geometric analysis to understand the working mechanism of deep learning from high-dimensional geometry is offered. Then, different forms of generative models like GAN, VAE, normalizing flows, optimal transport, and so on are described from a unified geometric perspective, showing that they actually come from statistical distance-minimization problems. Because this book contains up-to-date information from both a practical and theoretical point of view, it can be used as an advanced deep learning textbook in universities or as a reference source for researchers interested in

acquiring the latest deep learning algorithms and their underlying principles. In addition, the book has been prepared for a codeshare course for both engineering and mathematics students, thus much of the content is interdisciplinary and will appeal to students from both disciplines.

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

Convolutional Neural Networks for Medical Image Processing Applications - Saban Ozturk 2022-12-23

The rise in living standards increases the expectation of people in almost every field. At the forefront is health. Over the past few

centuries, there have been major developments in healthcare. Medical device technology and developments in artificial intelligence (AI) are among the most important ones. The improving technology and our ability to harness the technology effectively by means such as AI have led to unprecedented advances, resulting in early diagnosis of diseases. AI algorithms enable the fast and early evaluation of images from medical devices to maximize the benefits. While developments in the field of AI were quickly adapted to the field of health, in some cases this contributed to the formation of innovative artificial intelligence algorithms. Today, the most effective artificial intelligence method is accepted as deep learning. Convolutional neural network (CNN) architectures are deep learning algorithms used for image processing. This book contains applications of CNN methods. The content is quite extensive, including the application of different CNN methods to various medical image processing problems. Readers

will be able to analyze the effects of CNN methods presented in the book in medical applications.

Elements of Deep Learning for Computer Vision - Bharat Sikka 2021-06-24

Conceptualizing deep learning in computer vision applications using PyTorch and Python libraries. **KEY FEATURES** ● Covers a variety of computer vision projects, including face recognition and object recognition such as Yolo, Faster R-CNN. ● Includes graphical representations and illustrations of neural networks and teaches how to program them. ● Includes deep learning techniques and architectures introduced by Microsoft, Google, and the University of Oxford. **DESCRIPTION** Elements of Deep Learning for Computer Vision gives a thorough understanding of deep learning and provides highly accurate computer vision solutions while using libraries like PyTorch. This book introduces you to Deep Learning and explains all the concepts required to understand

the basic working, development, and tuning of a neural network using Pytorch. The book then addresses the field of computer vision using two libraries, including the Python wrapper/version of OpenCV and PIL. After establishing and understanding both the primary concepts, the book addresses them together by explaining Convolutional Neural Networks(CNNs). CNNs are further elaborated using top industry standards and research to explain how they provide complicated Object Detection in images and videos, while also explaining their evaluation. Towards the end, the book explains how to develop a fully functional object detection model, including its deployment over APIs. By the end of this book, you are well-equipped with the role of deep learning in the field of computer vision along with a guided process to design deep learning solutions. **WHAT YOU WILL LEARN** ● Get to know the mechanism of deep learning and how neural networks operate. ● Learn to develop a highly accurate neural

network model. ● Access to rich Python libraries to address computer vision challenges. ● Build deep learning models using PyTorch and learn how to deploy using the API. ● Learn to develop Object Detection and Face Recognition models along with their deployment. WHO THIS BOOK IS FOR This book is for the readers who aspire to gain a strong fundamental understanding of how to infuse deep learning into computer vision and image processing applications. Readers are expected to have intermediate Python skills. No previous knowledge of PyTorch and Computer Vision is required. TABLE OF CONTENTS 1. An Introduction to Deep Learning 2. Supervised Learning 3. Gradient Descent 4. OpenCV with Python 5. Python Imaging Library and Pillow 6. Introduction to Convolutional Neural Networks 7. GoogLeNet, VGGNet, and ResNet 8. Understanding Object Detection 9. Popular Algorithms for Object Detection 10. Faster RCNN with PyTorch and YoloV4 with Darknet 11. Comparing Algorithms and API Deployment

with Flask 12. Applications in Real World
Deep Learning and Convolutional Neural Networks for Medical Imaging and Clinical Informatics Le Lu 2019-09-19

This book reviews the state of the art in deep learning approaches to high-performance robust disease detection, robust and accurate organ segmentation in medical image computing (radiological and pathological imaging modalities), and the construction and mining of large-scale radiology databases. It particularly focuses on the application of convolutional neural networks, and on recurrent neural networks like LSTM, using numerous practical examples to complement the theory. The book's chief features are as follows: It highlights how deep neural networks can be used to address new questions and protocols, and to tackle current challenges in medical image computing; presents a comprehensive review of the latest research and literature; and describes a range of different methods that employ deep learning for

object or landmark detection tasks in 2D and 3D medical imaging. In addition, the book examines a broad selection of techniques for semantic segmentation using deep learning principles in medical imaging; introduces a novel approach to text and image deep embedding for a large-scale chest x-ray image database; and discusses how deep learning relational graphs can be used to organize a sizable collection of radiology findings from real clinical practice, allowing semantic similarity-based retrieval. The intended reader of this edited book is a professional engineer, scientist or a graduate student who is able to comprehend general concepts of image processing, computer vision and medical image analysis. They can apply computer science and mathematical principles into problem solving practices. It may be necessary to have a certain level of familiarity with a number of more advanced subjects: image formation and enhancement, image understanding, visual recognition in medical applications, statistical

learning, deep neural networks, structured prediction and image segmentation.

Neural Networks and Deep Learning - Pat Nakamoto 2018-06-30

Ready to crank up a deep neural network to get your self-driving car pick up the kids from school? Want to add 'Neural Networks' and '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! Check out this exceptional bundle of 3 books... This bundle contains 3 books: Book 1: Neural Networks & Deep Learning: Deep Learning explained to your granny - A visual introduction for beginners who want to make their own Deep Learning Neural Network... 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! * 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

Book2: Deep Learning: Deep Learning explained to your granny - A guide for Beginners... 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 Book 3: Blockchain Blueprint: The ultimate guide to understanding blockchain, cryptocurrencies, smart contracts and the future of money The current emerging innovation of this decade may be the connected world of computing relying on blockchain encryption. The attention given to this technology by global giant players suggests that it will become the operational philosophy of the economic system of the future, ranging across many industries. Blockchain can become the solution we needed

for speeding up the economy and transactions in order to keep up with our multi-device connected world. In this book, high tech expert Pat Nakamoto answers your questions concerning the future of Blockchain technology along with addressing different major developments linked to it, like Smart Contracts, Fintech and Ethereum. Hit download. Now!

Guide to Convolutional Neural Networks -

Hamed Habibi Aghdam 2017-05-17

This must-read text/reference introduces the fundamental concepts of convolutional neural networks (ConvNets), offering practical guidance on using libraries to implement ConvNets in applications of traffic sign detection and classification. The work presents techniques for optimizing the computational efficiency of ConvNets, as well as visualization techniques to better understand the underlying processes. The proposed models are also thoroughly evaluated from different perspectives, using exploratory and quantitative analysis. Topics and features:

explains the fundamental concepts behind training linear classifiers and feature learning; discusses the wide range of loss functions for training binary and multi-class classifiers; illustrates how to derive ConvNets from fully connected neural networks, and reviews different techniques for evaluating neural networks; presents a practical library for implementing ConvNets, explaining how to use a Python interface for the library to create and assess neural networks; describes two real-world examples of the detection and classification of traffic signs using deep learning methods; examines a range of varied techniques for visualizing neural networks, using a Python interface; provides self-study exercises at the end of each chapter, in addition to a helpful glossary, with relevant Python scripts supplied at an associated website. This self-contained guide will benefit those who seek to both understand the theory behind deep learning, and to gain hands-on experience in implementing

ConvNets in practice. As no prior background knowledge in the field is required to follow the material, the book is ideal for all students of computer vision and machine learning, and will also be of great interest to practitioners working on autonomous cars and advanced driver assistance systems.

Deep Learning for Vision Systems - Mohamed Elgendy 2020-11-10

How does the computer learn to understand what it sees? Deep Learning for Vision Systems answers that by applying deep learning to computer vision. Using only high school algebra, this book illuminates the concepts behind visual intuition. You'll understand how to use deep learning architectures to build vision system applications for image generation and facial recognition. Summary Computer vision is central to many leading-edge innovations, including self-driving cars, drones, augmented reality, facial recognition, and much, much more. Amazing new computer vision applications are developed

every day, thanks to rapid advances in AI and deep learning (DL). Deep Learning for Vision Systems teaches you the concepts and tools for building intelligent, scalable computer vision systems that can identify and react to objects in images, videos, and real life. With author Mohamed Elgendy's expert instruction and illustration of real-world projects, you'll finally grok state-of-the-art deep learning techniques, so you can build, contribute to, and lead in the exciting realm of computer vision! Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the technology How much has computer vision advanced? One ride in a Tesla is the only answer you'll need. Deep learning techniques have led to exciting breakthroughs in facial recognition, interactive simulations, and medical imaging, but nothing beats seeing a car respond to real-world stimuli while speeding down the highway. About the book How does the computer learn to

understand what it sees? Deep Learning for Vision Systems answers that by applying deep learning to computer vision. Using only high school algebra, this book illuminates the concepts behind visual intuition. You'll understand how to use deep learning architectures to build vision system applications for image generation and facial recognition. What's inside Image classification and object detection Advanced deep learning architectures Transfer learning and generative adversarial networks DeepDream and neural style transfer Visual embeddings and image search About the reader For intermediate Python programmers. About the author Mohamed Elgendy is the VP of Engineering at Rakuten. A seasoned AI expert, he has previously built and managed AI products at Amazon and Twilio. Table of Contents PART 1 - DEEP LEARNING FOUNDATION 1 Welcome to computer vision 2 Deep learning and neural networks 3 Convolutional neural networks 4 Structuring DL projects and hyperparameter

tuning PART 2 - IMAGE CLASSIFICATION AND DETECTION 5 Advanced CNN architectures 6 Transfer learning 7 Object detection with R-CNN, SSD, and YOLO PART 3 - GENERATIVE MODELS AND VISUAL EMBEDDINGS 8 Generative adversarial networks (GANs) 9 DeepDream and neural style transfer 10 Visual embeddings

Applied Deep Learning with Pytorch - Hyatt Saleh 2019-04-26

Implement techniques such as image classification and natural language processing (NLP) by understanding the different neural network architectures Key Features Understand deep learning and how it can solve complex real-world problems Apply deep learning for image classification and text processing using neural networks Develop deep learning solutions for tasks such as basic classification and solving style transfer problems Book Description Machine learning is rapidly becoming the most preferred way of solving data problems, thanks

to the huge variety of mathematical algorithms that find patterns, which are otherwise invisible to us. Applied Deep Learning with PyTorch takes your understanding of deep learning, its algorithms, and its applications to a higher level. The book begins by helping you browse through the basics of deep learning and PyTorch. Once you are well versed with the PyTorch syntax and capable of building a single-layer neural network, you will gradually learn to tackle more complex data problems by configuring and training a convolutional neural network (CNN) to perform image classification. As you progress through the chapters, you'll discover how you can solve an NLP problem by implementing a recurrent neural network (RNN). By the end of this book, you'll be able to apply the skills and confidence you've gathered along your learning process to use PyTorch for building deep learning solutions that can solve your business data problems. What you will learn Detect a variety of data problems to which you can apply

deep learning solutions Learn the PyTorch syntax and build a single-layer neural network with it Build a deep neural network to solve a classification problem Develop a style transfer model Implement data augmentation and retrain your model Build a system for text processing using a recurrent neural network Who this book is for Applied Deep Learning with PyTorch is designed for data scientists, data analysts, and developers who want to work with data using deep learning techniques. Anyone looking to explore and implement advanced algorithms with PyTorch will also find this book useful. Some working knowledge of Python and familiarity with the basics of machine learning are a must. However, knowledge of NumPy and pandas will be beneficial, but not essential. **Math for Deep Learning** - Ronald T. Kneusel 2021-12-07 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.

Supervised Machine Learning for Text Analysis in R - Emil Hvitfeldt 2021-10-22

Text data is important for many domains, from healthcare to marketing to the digital humanities, but specialized approaches are necessary to create features for machine

learning from language. *Supervised Machine Learning for Text Analysis in R* explains how to preprocess text data for modeling, train models, and evaluate model performance using tools from the tidyverse and tidymodels ecosystem. Models like these can be used to make predictions for new observations, to understand what natural language features or characteristics contribute to differences in the output, and more. If you are already familiar with the basics of predictive modeling, use the comprehensive, detailed examples in this book to extend your skills to the domain of natural language processing. This book provides practical guidance and directly applicable knowledge for data scientists and analysts who want to integrate unstructured text data into their modeling pipelines. Learn how to use text data for both regression and classification tasks, and how to apply more straightforward algorithms like regularized regression or support vector machines as well as deep

learning approaches. Natural language must be dramatically transformed to be ready for computation, so we explore typical text preprocessing and feature engineering steps like tokenization and word embeddings from the ground up. These steps influence model results in ways we can measure, both in terms of model metrics and other tangible consequences such as how fair or appropriate model results are.

Fundamentals of Deep Learning and Computer Vision - Nikhil Singh 2020-02-24

Master Computer Vision concepts using Deep Learning with easy-to-follow steps
DESCRIPTION This book starts with setting up a Python virtual environment with the deep learning framework TensorFlow and then introduces the fundamental concepts of TensorFlow. Before moving on to Computer Vision, you will learn about neural networks and related aspects such as loss functions, gradient descent optimization, activation functions and how backpropagation works for training multi-

layer perceptrons. To understand how the Convolutional Neural Network (CNN) is used for computer vision problems, you need to learn about the basic convolution operation. You will learn how CNN is different from a multi-layer perceptron along with a thorough discussion on the different building blocks of the CNN architecture such as kernel size, stride, padding, and pooling and finally learn how to build a small CNN model. Next, you will learn about different popular CNN architectures such as AlexNet, VGGNet, Inception, and ResNets along with different object detection algorithms such as RCNN, SSD, and YOLO. The book concludes with a chapter on sequential models where you will learn about RNN, GRU, and LSTMs and their architectures and understand their applications in machine translation, image/video captioning and video classification. KEY FEATURES Setting up the Python and TensorFlow environment Learn core Tensorflow concepts with the latest TF version 2.0 Learn

Deep Learning for computer vision applications
Understand different computer vision concepts
and use-cases Understand different state-of-the-
art CNN architectures Build deep neural
networks with transfer Learning using features
from pre-trained CNN models Apply computer
vision concepts with easy-to-follow code in
Jupyter Notebook WHAT WILL YOU LEARN This
book will help the readers to understand and
apply the latest Deep Learning technologies to
different interesting computer vision
applications without any prior domain
knowledge of image processing. Thus, helping
the users to acquire new skills specific to
Computer Vision and Deep Learning and build
solutions to real-life problems such as Image
Classification and Object Detection. This book
will serve as a basic guide for all the beginners
to master Deep Learning and Computer Vision
with lucid and intuitive explanations using basic
mathematical concepts. It also explores these
concepts with popular the deep learning

framework TensorFlow. WHO THIS BOOK IS
FOR This book is for all the Data Science
enthusiasts and practitioners who intend to
learn and master Computer Vision concepts and
their applications using Deep Learning. This
book assumes a basic Python understanding with
hands-on experience. A basic senior secondary
level understanding of Mathematics will help the
reader to make the best out of this book. Table
of Contents 1. Introduction to TensorFlow 2.
Introduction to Neural Networks 3.
Convolutional Neural Network 4. CNN
Architectures 5. Sequential Models
Computer Vision -- ECCV 2014 David Fleet
2014-09-22

The seven-volume set comprising LNCS volumes
8689-8695 constitutes the refereed proceedings
of the 13th European Conference on Computer
Vision, ECCV 2014, held in Zurich, Switzerland,
in September 2014. The 363 revised papers
presented were carefully reviewed and selected
from 1444 submissions. The papers are

organized in topical sections on tracking and activity recognition; recognition; learning and inference; structure from motion and feature matching; computational photography and low-level vision; vision; segmentation and saliency; context and 3D scenes; motion and 3D scene analysis; and poster sessions.

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.

Neural Networks and Statistical Learning

Lin Du 2019-09-12

This book provides a broad yet detailed introduction to neural networks and machine learning in a statistical framework. A single,

comprehensive resource for study and further research, it explores the major popular neural network models and statistical learning approaches with examples and exercises and allows readers to gain a practical working understanding of the content. This updated new edition presents recently published results and includes six new chapters that correspond to the recent advances in computational learning theory, sparse coding, deep learning, big data and cloud computing. Each chapter features state-of-the-art descriptions and significant research findings. The topics covered include: • multilayer perceptron; • the Hopfield network; • associative memory models; • clustering models and algorithms; • the radial basis function network; • recurrent neural networks; • nonnegative matrix factorization; • independent component analysis; • probabilistic and Bayesian networks; and • fuzzy sets and logic. Focusing on the prominent accomplishments and their practical aspects, this book provides academic

and technical staff, as well as graduate students and researchers with a solid foundation and comprehensive reference on the fields of neural networks, pattern recognition, signal processing, and machine learning.

Understanding Deep Learning - Chitta Ranjan
2020-12-24

Think of deep learning as an art of cooking. One way to cook is to follow a recipe. But when we learn how the food, the spices, and the fire behave, we make our creation. And an understanding of the "how" transcends the creation. Likewise, an understanding of the "how" transcends deep learning. In this spirit, this book presents the deep learning constructs, their fundamentals, and how they behave. Baseline models are developed alongside, and concepts to improve them are exemplified. Topics covered in the book include:- Multilayer Perceptrons- Long- and short-term Memory Networks- Convolutional Neural Networks- Autoencoders Every topic is thoroughly explained

and illustrated graphically. Moreover, implementations in TensorFlow are given for developing a complete understanding.

Practical Convolutional Neural Networks -
Mohit Sewak 2018-02-27

This book helps you master CNN, from the basics to the most advanced concepts in CNN such as GANs, instance classification and attention mechanism for vision models and more. You will implement advanced CNN models using complex image and video datasets. By the end of the book you will learn CNN's best practices to implement smart ConvNet ...

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.

Data Science - Gyanendra K. Verma 2021-08-19
This book targets an audience with a basic understanding of deep learning, its architectures, and its application in the multimedia domain. Background in machine learning is helpful in exploring various aspects of deep learning. Deep learning models have a major impact on multimedia research and raised the performance bar substantially in many of the standard evaluations. Moreover, new multi-modal challenges are tackled, which older

systems would not have been able to handle. However, it is very difficult to comprehend, let alone guide, the process of learning in deep neural networks, there is an air of uncertainty about exactly what and how these networks learn. By the end of the book, the readers will have an understanding of different deep learning approaches, models, pre-trained models, and familiarity with the implementation of various deep learning algorithms using various frameworks and libraries.

Insi de Deep Learni ng Edward Raff 2022-07-05

Journey through the theory and practice of modern deep learning, and apply innovative techniques to solve everyday data problems. In *Inside Deep Learning*, you will learn how to:

- Implement deep learning with PyTorch
- Select the right deep learning components
- Train and evaluate a deep learning model
- Fine tune deep learning models to maximize performance
- Understand deep learning terminology
- Adapt existing PyTorch code to solve new problems

Inside Deep Learning is an accessible guide to implementing deep learning with the PyTorch framework. It demystifies complex deep learning concepts and teaches you to understand the vocabulary of deep learning so you can keep pace in a rapidly evolving field. No detail is skipped—you'll dive into math, theory, and practical applications. Everything is clearly explained in plain English. About the technology Deep learning doesn't have to be a black box! Knowing how your models and algorithms actually work gives you greater control over your results. And you don't have to be a mathematics expert or a senior data scientist to grasp what's going on inside a deep learning system. This book gives you the practical insight you need to understand and explain your work with confidence. About the book *Inside Deep Learning* illuminates the inner workings of deep learning algorithms in a way that even machine learning novices can understand. You'll explore deep learning concepts and tools through plain

language explanations, annotated code, and dozens of instantly useful PyTorch examples. Each type of neural network is clearly presented without complex math, and every solution in this book can run using readily available GPU hardware! What's inside

Select the right deep learning components
Train and evaluate a deep learning model
Fine tune deep learning models to maximize performance
Understand deep learning terminology

About the reader
For Python programmers with basic machine learning skills.

About the author
Edward Raff is a Chief Scientist at Booz Allen Hamilton, and the

author of the JSAT machine learning library.

Table of Contents

PART 1 FOUNDATIONAL METHODS

1 The mechanics of learning
2 Fully connected networks
3 Convolutional neural networks
4 Recurrent neural networks
5 Modern training techniques
6 Common design building blocks

PART 2 BUILDING ADVANCED NETWORKS

7 Autoencoding and self-supervision
8 Object detection
9 Generative adversarial networks
10 Attention mechanisms
11 Sequence-to-sequence
12 Network design alternatives to RNNs
13 Transfer learning
14 Advanced building blocks