Understanding Deep Convolutional Neural Networks With A

In recent years, deep learning models have been widely used and are behind major breakthroughs across many fields. Deep learning models are usually considered to be black boxes due to their large model structures and complicated hierarchical nonlinear transformations. As deep learning technology continues to develop, the understanding of deep learning models is raising concerns, such as the understanding of the training and prediction behaviors and the internal mechanism of models. In this thesis, we study the model understanding problem of deep neural networks from the perspective of piecewise linear property. First, we introduce the piecewise linear property. Next, we review the role and progress of deep learning understanding from the perspective of the piecewise linear property. The piecewise linear property reveals that deep neural networks with piecewise linear activation functions can generally divide the input space into a number of small disjointed regions that correspond to a local linear function within each region. Next, we investigate two typical understanding problems, namely model interpretation, and model complexity. In particular, we provide a series of derivations and analyses of the piecewise linear property of deep neural networks with piecewise linear activation functions. We propose an approach for interpreting the predictions given by such models based on the piecewise linear property. Next, we propose a method to provide local interpretation to a black box deep model by mimicking a piecewise linear approximation from the deep model. Then, we study deep neural networks with curve activation functions with the aim of providing piecewise linear approximations for these networks that would let them benefit from the piecewise linear property. After proposing a piecewise linear approximation framework, we investigate model complexity and model interpretation using the approximation. The thesis concludes by discussing future directions for understanding deep neural networks from the perspective of the piecewise linear property.

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.

Neural Networks and their implementation decoded with TensorFlow About This Book Develop a strong background in neural network programming from scratch, using the popular Tensorflow library. Use Tensorflow to implement different kinds of neural networks – from simple feedforward neural networks to multilayered perceptrons, CNNs, RNNs and more. A highly practical guide including real-world datasets and use-cases to simplify your understanding of neural networks and their implementation. Who This Book Is For This book is meant for developers with a statistical background who want to work with neural networks. Though we will be using TensorFlow as the underlying library for neural networks, book can be used as a generic resource to bridge the gap between the math and the implementation of deep learning. If you have some understanding of Tensorflow and Python and want to learn what happens at a level lower than the plain API syntax, this book is for you. What You Will Learn Learn Linear Algebra and mathematics behind neural network. Dive deep into Neural networks from the basic to advanced concepts like CNN, RNN Deep Belief Networks, Deep Feedforward Networks. Explore Optimization techniques for solving problems like Local minima, Global minima, Saddle points Learn through real world examples like Sentiment Analysis. Train different types of generative models and explore autoencoders. Explore TensorFlow as an example of deep learning implementation. In Detail If you're aware of the buzz surrounding the terms such as "machine learning," "artificial intelligence," or "deep learning," you might know what neural networks are. Ever wondered how they help in solving complex computational problem efficiently, or how to train efficient neural networks? This book will teach you just that. You will start by getting a quick overview of the popular TensorFlow library and how it is used to train different neural networks. You will get a thorough understanding of the fundamentals and basic math for neural networks and why TensorFlow is a popular choice Then, you will proceed to implement a simple feed forward neural network. Next you will master optimization techniques and algorithms for neural networks using TensorFlow. Further, you will learn to implement some more complex types of neural networks such as convolutional neural networks, recurrent neural networks, and Deep Belief Networks. In the course of the book, you will be working on real-world datasets to get a hands-on understanding of neural network programming. You will also get to train generative models and will learn the applications of autoencoders. By the end of this book, you will have a fair understanding of how you can leverage the power of TensorFlow to train neural networks of varying complexities, without any hassle. While you are learning about various neural network implementations you will learn the underlying mathematics and linear algebra and how they map to the appropriate TensorFlow constructs. Style and Approach This book is designed to give you just the right number of concepts to back up the examples. With real-world use cases and problems solved, this book is a handy guide for you. Each concept is backed by a generic and real-world problem, followed by a variation, making you independent and able to solve any problem with neural networks. All of the content is demystified by a simple and straightforward approach.

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

The Principles of Deep Learning Theory

Deep Learning

Image Classification, Object Detection, and Face Recognition in Python

Beginner's Guide To Convolutional Neural Networks In Python

Neural Networks and Deep Learning

Demystify Neural Networks with Pytorch

A Complete Guide to become an Expert in Deep Learning and Computer Vision

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

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.

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!

Deep Learning, a type of Artificial Intelligence, is transforming many industries including transportation, health care and mobile computing. The main actors behind deep learning are deep neural networks (DNNs). These artificial brains have demonstrated impressive performance on many challenging tasks such as synthesizing and recognizing speech, driving cars, and even detecting cancer from medical scans. Given their excellent performance and widespread applications in everyday life, it is important to understand: (1) how DNNs function internally; (2) why they perform so well; and (3) when they fail. Answering these questions would allow end-users (e.g. medical doctors harnessing deep learning to assist them in diagnosis) to gain deeper insights into how these models behave, and therefore more confidence in utilizing the technology in important real-world applications. Artificial neural networks traditionally had been treated as black boxes—little was known about how they arrive at a decision when an input is present. Similarly, in neuroscience, understanding how biological brains work has also been a long-standing quest. Neuroscientists have discovered neurons in human brains that selectively fire in response to specific, abstract concepts such as Halle Berry or Bill Clinton, informing the discussion of whether learned neural codes are local or distributed. These neurons were identified by finding the preferred stimuli (here, images) that highly excite a specific neuron, which was accomplished by showing subjects many different images while recording a target neuron's activation. Inspired by such neuroscience techniques, my Ph.D. study produced a series of visualization methods that synthesize the preferred stimuli for each neuron in DNNs to shed more light into (1) the weaknesses of DNNs, which raise serious concerns about their widespread deployment in critical sectors of our economy and society; and (2) how DNNs function internally. Some of the notable findings are summarized as follows. First, DNNs are easily fooled in that it is possible to produce images that are visually unrecognizable to humans, but that state-of-the-art DNNs classify as familiar objects with near certainty confidence (i.e. labeling white-noise images as "school bus"). These images can be optimized to fool the DNN regardless of whether we treat the network as a white- or black-box (i.e. we have access to the network parameters or not). These results shed more light into the inner workings of DNNs and also question the security and reliability of deep learning applications. Second, our visualization methods reveal that DNNs can automatically learn a hierarchy of increasingly abstract features from the input space that are useful to solve a given task. In addition, we also found that neurons in DNNs are often multifaceted in that a single neuron fires for a variety of different input patterns (i.e. it is invariant to changes in the input). These observations align with the common wisdom previously established for both human visual cortex and DNNs. Lastly, many machine learning hobbyists and scientists have successfully applied our methods to visualize their own DNNs or even generate high-quality art images. We also turn the visualization frameworks into (1) an art generator algorithm, and (2) a state-of-the-art image generative model, making contributions to the fields of evolutionary computation and generative modeling, respectively.

A Textbook

Visualizing and Understanding Deep Neural Networks

Making AI Less Susceptible to Adversarial Trickery

Convolutional Neural Networks In Python

2 Manuscripts - Deep Learning With Keras And Convolutional Neural Networks In Python

Learn the skills you need to develop your own next-generation deep learning models with TensorFlow and Keras

Principles and Labs for Deep Learning provides the knowledge and techniques needed to help readers design and develop deep learning models. Deep Learning techniques are introduced through theory, comprehensively illustrated, explained through the TensorFlow source code examples, and analyzed through the visualization of results. The structured methods and labs provided by Dr. Huang and Dr. Le enable readers to become proficient in TensorFlow to build deep Convolutional Neural Networks (CNNs) through custom APIs, high-level Keras APIs, Keras Applications, and TensorFlow Hub. Each chapter has one corresponding Lab with step-by-step instruction to help the reader practice and accomplish a specific learning outcome. Deep Learning has been successfully applied in diverse fields such as computer vision, audio processing, robotics, natural language processing, bioinformatics and chemistry. Because of the huge scope of knowledge in Deep Learning, a lot of time is required to understand and deploy useful, working applications, hence the importance of this new resource. Both theory lessons and experiments are included in each chapter to introduce the techniques and provide source code examples to practice using them. All Labs for this book are placed on GitHub to facilitate the download. The book is written based on the assumption that the reader knows basic Python for programming and basic Machine Learning. Introduces readers to the usefulness of neural networks and Deep Learning methods Provides readers with in-depth understanding of the architecture and operation of Deep Convolutional Neural Networks Demonstrates the visualization needed for designing neural networks Provides readers with an in-depth understanding of regression problems, binary classification problems, multi-category classification problems, Variational Auto-Encoder, Generative Adversarial Network, and Object detection

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

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 NOW and learn more about Convolutional Neural Networks in Python!

Build, implement and scale distributed deep learning models for large-scale datasets About This Book Get to grips with the deep learning concepts and set up Hadoop to put them to use Implement and parallelize deep learning models on Hadoop's YARN framework A comprehensive tutorial to distributed deep learning with Hadoop Who This Book Is For If you are a data scientist who wants to learn how to perform deep learning on Hadoop, this is the book for you. Knowledge of the basic machine learning concepts and some understanding of Hadoop is required to make the best use of this book. What You Will Learn Explore Deep Learning and various models associated with it Understand the challenges of implementing distributed deep learning with Hadoop and how to overcome it Implement Convolutional Neural Network (CNN) with deeplearning4j Delve into the implementation of Restricted Boltzmann Machines (RBM) Understand the mathematical explanation for implementing Recurrent Neural Networks (RNN) Get hands on practice of deep learning and their implementation with Hadoop. In Detail This book will teach you how to deploy large-scale dataset in deep neural networks with Hadoop for optimal performance. Starting with understanding what deep learning is, and what the various models associated with deep neural networks are, this book will then show you how to set up the Hadoop environment for deep learning. In this book, you will also learn how to overcome the challenges that you face while implementing distributed deep learning with large-scale unstructured datasets. The book will also show you how you can implement and parallelize the widely used deep learning models such as Deep Belief Networks, Convolutional Neural Networks, Recurrent Neural Networks, Restricted Boltzmann Machines and autoencoder using the popular deep learning library deeplearning4j. Get in-depth mathematical explanations and visual representations to help you understand the design and implementations of Recurrent Neural network and Denoising AutoEncoders with deeplearning4j. To give you a more practical perspective, the book will also teach you the implementation of large-scale video processing, image processing and natural language processing on Hadoop. By the end of this book, you will know how to deploy various deep neural networks in distributed systems using Hadoop. Style and approach This book takes a comprehensive, step-by-step approach to implement efficient deep learning models on Hadoop. It starts from the basics and builds the readers' knowledge as they strengthen their understanding of the concepts. Practical examples are included in every step of the way to supplement the theory.

A Practitioner's Approach

A Signal Processing Perspective

Efficient Processing of Deep Neural Networks

Principles and Labs for Deep Learning

Geometry of Deep Learning

Inside Deep Learning

A Practical Application to Traffic-Sign Detection and Classification

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.

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.

As deep neural networks (DNNs) become increasingly common in real-world applications, the potential to deliberately "fool" them with data that wouldn’t trick a human presents a new attack vector. This practical book examines real-world scenarios where DNNs—the algorithms intrinsic to much of AI—are used daily to process image, audio, and video data. Author Katy Warr considers attack motivations, the risks posed by this adversarial input, and methods for increasing AI robustness to these attacks. If you’re a data scientist developing DNN algorithms, a security architect interested in how to make AI systems more resilient to attack, or someone fascinated by the differences between artificial and biological perception, this book is for you. Delve into DNNs and discover how they could be tricked by adversarial input Investigate methods used to generate adversarial input capable of fooling DNNs Explore real-world scenarios and model the adversarial threat Evaluate neural network robustness; learn methods to increase resilience of AI systems to adversarial data Examine some ways in which AI might become better at mimicking human perception in years to come

Take a hands-on approach to understanding deep learning and build smart applications that can recognize images and interpret text Key Features Understand how to implement deep learning with TensorFlow and Keras Learn the fundamentals of computer vision and image recognition Study the architecture of different neural networks Book Description Are you fascinated by how deep learning powers intelligent applications such as self-driving cars, virtual assistants, facial recognition devices, and chatbots to process data and solve complex problems? Whether you are familiar with machine learning or are new to this domain, The Deep Learning Workshop will make it easy for you to understand deep learning with the help of interesting examples and exercises throughout. The book starts by highlighting the relationship between deep learning, machine learning, and artificial intelligence and helps you get comfortable with the TensorFlow 2.0 programming structure using hands-on exercises. You'll understand neural networks, the structure of a perceptron, and how to use TensorFlow to create and train models. The book will then let you explore the fundamentals of computer vision by performing image recognition exercises with convolutional neural networks (CNNs) using Keras. As you advance, you'll be able to make your model more powerful by implementing text embedding and sequencing the data using popular deep learning solutions. Finally, you'll get to grips with bidirectional recurrent neural networks (RNNs) and build generative adversarial networks (GANs) for image synthesis. By the end of this deep learning book, you'll have learned the skills essential for building deep learning models with TensorFlow and Keras. What you will learn Understand how deep learning, machine learning, and artificial intelligence are different Develop multilayer deep neural networks with TensorFlow Implement deep neural networks for multiclass classification using Keras Train CNN models for image recognition Handle sequence data and use it in conjunction with RNNs Build a GAN to generate high-quality synthesized images Who this book is for If you are interested in machine learning and want to create and train deep learning models using TensorFlow and Keras, this workshop is for you. A solid understanding of Python and its packages, along with basic machine learning concepts, will help you to learn the topics quickly.

Become an expert at designing, building, and improving advanced neural network models using R

Applied Deep Learning

The Ultimate Guide to Understand Deep Neural Networks with Python Through PyTorch, TensorFlow and Keras. Discover the Ethical Implications of Deep Learning in the New World

Build and deploy distributed deep learning applications on Apache Spark

Advanced Applied Deep Learning

Deep Learning for Vision Systems

A Case-Based Approach to Understanding Deep Neural Networks

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!

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.

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

Speed up the design and implementation of deep learning solutions using Apache Spark Key FeaturesExplore the world of distributed deep learning with Apache SparkTrain neural networks with deep learning libraries such as BigDL and TensorFlowDevelop Spark deep learning applications to intelligently handle large and complex datasetsBook Description Deep learning is a subset of machine learning where datasets with several layers of complexity can be processed. Hands-On Deep Learning with Apache Spark addresses the sheer complexity of technical and analytical parts and the speed at which deep learning solutions can be implemented on Apache Spark. The book starts with the fundamentals of Apache Spark and deep learning. You will set up Spark for deep learning, learn principles of distributed modeling, and understand different types of neural nets. You will then implement deep learning models, such as convolutional neural networks (CNNs), recurrent neural networks (RNNs), and long short-term memory (LSTM) on Spark. As you progress through the book, you will gain hands-on experience of what it takes to understand the complex datasets you are dealing with. During the course of this book, you will use popular deep learning frameworks, such as TensorFlow, Deeplearning4j, and Keras to train your distributed models. By the end of this book, you'll have gained experience with the implementation of your models on a variety of use cases. What you will learnUnderstand the basics of deep learningSet up Apache Spark for deep learningUnderstand the principles of distribution modeling and different types of neural networksObtain an understanding of deep learning algorithmsDiscover textual analysis and deep learning with SparkUse popular deep learning frameworks, such as Deeplearning4j, TensorFlow, and KerasExplore popular deep learning algorithms Who this book is for If you are a Scala developer, data scientist, or data analyst who wants to learn how to use Spark for implementing efficient deep learning models, Hands-On Deep Learning with Apache Spark is for you. Knowledge of the core machine learning concepts and some exposure to Spark will be helpful.

Advanced Deep Learning with R

Deep Learning for Computer Vision

Unleash the power of TensorFlow to train efficient neural networks

The Deep Learning Workshop

Math for Deep Learning

Explainable AI: Interpreting, Explaining and Visualizing Deep Learning

Strengthening Deep Neural Networks

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.

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 workSave weights and models on diskPause training and restart it at a later stage Use hardware acceleration (GPUs) in your codeWork with the Dataset TensorFlow abstraction and use pre-trained models and transfer learningRemove and add layers to pre-trained networks to adapt them to your specific projectApply 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.

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 LearningA Case-Based Approach to Understanding Deep Neural NetworksApress

Elements of Deep Learning for Computer Vision

Robustness, Uncertainty Quantification, and Insights Towards Safety

Deep Learning from Scratch

Interpretable Machine Learning

Neural Network Programming with TensorFlow

Deep Neural Networks and Data for Automated Driving

Guide to Convolutional Neural Networks

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.

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

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.

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.

Neural Networks and Deep Learning, Deep Learning, Blockchain Blueprint

Deep Learning with Hadoop

What You Need to Know to Understand Neural Networks

AI Neuroscience

Implement advanced deep learning models using Python

Understanding Deep Neural Networks from the Perspective of Piecewise Linear Property

Practical Convolutional Neural Networks

Gain expertise in advanced deep learning domains such as neural networks, meta-learning, graph neural networks, and memory augmented neural networks using the Python ecosystem Key Features Get to grips with building faster and more robust deep learning architectures Investigate and train convolutional neural network (CNN) models with GPU-accelerated libraries such as TensorFlow and PyTorch Apply deep neural networks (DNNs) to computer vision problems, NLP, and GANs Book Description In order to build robust deep learning systems, you'll need to understand everything from how neural networks work to training CNN models. In this book, you'll discover newly developed deep learning models, methodologies used in the domain, and their implementation based on areas of application. You'll start by understanding the building blocks and the math behind neural networks, and then move on to CNNs and their advanced applications in computer vision. You'll also learn to apply the most popular CNN architectures in object detection and image segmentation. Further on, you'll focus on variational autoencoders and GANs. You'll then use neural networks to extract sophisticated vector representations of words, before going on to cover various types of recurrent networks, such as LSTM and GRU. You'll even explore the attention mechanism to process sequential data without the help of recurrent neural networks (RNNs). Later, you'll use graph neural networks for processing structured data, along with covering meta-learning, which allows you to train neural networks with fewer training samples. Finally, you'll understand how to apply deep learning to autonomous vehicles. By the end of this book, you'll have mastered key deep learning concepts and the different applications of deep learning models in the real world. What you will learn Cover advanced and state-of-the-art neural network architectures Understand the theory and math behind neural networks Train DNNs and apply them to modern deep learning problems Use CNNs for object detection and image segmentation Implement generative adversarial networks (GANs) and variational autoencoders to generate new images Solve natural language processing (NLP) tasks, such as machine translation, using sequence-to-sequence models Understand DL techniques, such as meta-learning and graph neural networks Who this book is for This book is for data scientists, deep learning engineers and researchers, and AI developers who want to further their knowledge of deep learning and build innovative and unique deep learning projects. Anyone looking to get to grips with advanced use cases and methodologies adopted in the deep learning domain using real-world examples will also find this book useful. Basic understanding of deep learning concepts and working knowledge of the Python programming language is assumed.

Master Computer Vision concepts using Deep Learning with easy-to-follow steps Key Featuresa- Setting up the Python and TensorFlow environmenta- Learn core Tensorflow concepts with the latest TF version 2.0a- Learn Deep Learning for computer vision applications a- Understand different computer vision concepts and use-casesa- Understand different state-of-the-art CNN architectures a- Build deep neural networks with transfer Learning using features from pre-trained CNN modelsa- Apply computer vision concepts with easy-to-follow code in Jupyter NotebookDescriptionThis 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.What will you learnThis 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. Who this book is forThis 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 Contents1. Introduction to TensorFlow2. Introduction to Neural Networks 3. Convolutional Neural Network 4. CNN Architectures5. Sequential ModelsAbout the AuthorNikhil Singh is an accomplished data scientist and currently working as the Lead Data Scientist at Proarch IT Solutions Pvt. Ltd in London. He has experience in designing and delivering complex and innovative computer vision and NLP centred solutions for a large number of global companies. He has been an AI consultant to a few companies and mentored many apprentice Data Scientists. His LinkedIn Profile: https://www.linkedin.com/in/nikhil-singh-b953ba122/Paras Ahuja is a seasoned data science practitioner and currently working as the Lead Data Scientist at Reliance Jio in Hyderabad. He has good experience in designing and deploying deep learning-based Computer Vision and NLP-based solutions. He has experience in developing and implementing state-of-the-art automatic speech recognition systems.His LinkedIn Profile: https://www.linkedin.com/in/parasahuja

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.

Apply modern deep learning techniques to build and train deep neural networks using Gorgonia Key Features Gain a practical understanding of deep learning using Golang Build complex neural network models using Go libraries and Gorgonia Take your deep learning model from design to deployment with this handy guide Book Description Go is an open source programming language designed by Google for handling large-scale projects efficiently. The Go ecosystem comprises some really powerful deep learning tools such as DQN and CUDA. With this book, you'll be able to use these tools to train and deploy scalable deep learning models from scratch. This deep learning book begins by introducing you to a variety of tools and libraries available in Go. It then takes you through building neural networks, including activation functions and the learning algorithms that make neural networks tick. In addition to this, you'll learn how to build advanced architectures such as autoencoders, restricted Boltzmann machines (RBMs), convolutional neural networks (CNNs), recurrent neural networks (RNNs), and more. You'll also understand how you can scale model deployments on the AWS cloud infrastructure for training and inference. By the end of this book, you'll have mastered the art of building, training, and deploying deep learning models in Go to solve real-world problems. What you will learn Explore the Go ecosystem of libraries and communities for deep learning Get to grips with Neural Networks, their history, and how they work Design and implement Deep Neural Networks in Go Get a strong foundation of concepts such as Backpropagation and Momentum Build Variational Autoencoders and Restricted Boltzmann Machines using Go Build models with CUDA and benchmark CPU and GPU models Who this book is for This book is for data scientists, machine learning engineers, and AI developers who want to build state-of-the-art deep learning models using Go. Familiarity with basic machine learning concepts and Go programming is required to get the best out of this book.

Neural Networks and Deep Learning, Deep Learning, Big Data

Advanced Deep Learning with Python

Deep Learning with R

Deep Learning with Python

Deep Learning in Mining of Visual Content

Applied Deep Learning with Pytorch

Explore Deep Neural Network Architectures, PyTorch, Object Detection Algorithms, and Computer Vision Applications for Python Coders (English Edition)

This book doesn't have any superpowers or magic formula to help you master the art of neural networks and deep learning. We believe that such learning is all in your heart. You need to learn a concept by heart and then brainstorm its different possibilities. I don't claim that after reading this book you will become an expert in Python and Deep Learning Neural Networks. Instead, you will, for sure, have a basic understanding of deep learning and its implications and real-life applications. Most of the time, what confuses us is the application of a certain thing in our lives. Once we know that, we can relate the subject to that particular thing and learn. An interesting thing is that neural networks also learn the same way. This makes it easier to learn about them when we know the basics. Let's take a look at what this book has to offer: ● The basics of Python including data types, operators and numbers. ● Advanced programming in Python with Python expressions, types and much more. ● A comprehensive overview of deep learning and its link to the smart systems that we are now building. ● An overview of how artificial neural networks work in real life. ● An overview of PyTorch. ● An overview of TensorFlow. ● An overview of Keras. ● How to create a convolutional neural network. ● A comprehensive understanding of deep learning applications and its ethical implications, including in the present and future. This book offers you the basic knowledge about Python and Deep Learning Neural Networks that you will need to lay the foundation for future studies. This book will start you on the road to mastering the art of deep learning neural networks. When I say that I don't have the magic formula to make you learn, I mean it. My point is that you should learn Python coding and Python libraries to build neural networks by practicing hard. The more you practice, the better it is for your skills. It is only after thorough and in depth practice that you will be able to create your own programs. Unlike other books, I don't claim that this book will make you a master of deep learning after a single read. That's not realistic, in fact, it's even a bit absurd. What I claim is that you will definitely learn about the basics. The rest is practice. The more you practice the better you code.

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

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.

Step-by-step tutorials on deep learning neural networks for computer vision in python with Keras.

A practical guide to building and implementing neural network models using Go

Fundamentals of Deep Learning and Computer Vision

Hands-On Deep Learning with Apache Spark

A Guide to Convolutional Neural Networks for Computer Vision

Design and implement advanced next-generation AI solutions using TensorFlow and PyTorch

Math, Algorithms, Models

Hands-On Deep Learning with Go

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.

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.

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.

This volume develops an effective theory approach to understanding deep neural networks of practical relevance.

Building with Python from First Principles

Convolutional Neural Networks and Object Detection