

An Introduction To Convolutional Neural Networks

Thank you for downloading **An Introduction To Convolutional Neural Networks** . Maybe you have knowledge that, people have look hundreds times for their favorite novels like this An Introduction To Convolutional Neural Networks , but end up in malicious downloads.

Rather than enjoying a good book with a cup of coffee in the afternoon, instead they juggled with some harmful bugs inside their computer.

An Introduction To Convolutional Neural Networks is available in our digital library an online access to it is set as public so you can get it instantly.

Our book servers hosts in multiple locations, allowing you to get the most less latency time to download any of our books like this one.

Merely said, the An Introduction To Convolutional Neural Networks is universally compatible with any devices to read

Introduction to Deep Learning and Neural Networks with PythonTM - Ahmed Fawzy Gad 2020-12-07

Introduction to Deep Learning and Neural Networks with PythonTM: A Practical Guide is an intensive step-by-step guide for neuroscientists to fully understand, practice, and build neural networks. Providing math and PythonTM code examples to clarify neural network calculations, by book's end readers will fully understand how neural networks work starting from the simplest model $Y=X$ and building from scratch. Details and explanations are provided on how a generic gradient descent algorithm works based on mathematical and PythonTM examples, teaching you how to use the gradient descent algorithm to manually perform all calculations in both the forward and backward passes of training a neural network. Examines the practical side of deep learning and neural networks Provides a problem-based approach to building artificial neural networks using real data Describes PythonTM functions and features for neuroscientists Uses a careful tutorial approach to describe implementation of neural networks in PythonTM Features math and code examples (via companion website) with helpful instructions for easy implementation

Convolutional Neural Networks In Python - Frank Millstein 2020-07-06

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!

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.

Data-Driven Science and Engineering - Steven L. Brunton 2022-05-05

A textbook covering data-science and machine learning methods for modelling and control in engineering and science, with Python and MATLAB®.

Grokking Machine Learning - Luis Serrano 2021-12-14
Grokking Machine Learning presents machine learning algorithms and techniques in a way that anyone can understand. This book skips the confused academic jargon and offers clear explanations that require only basic algebra. As you go, you'll build interesting projects with Python, including models for spam detection and image recognition. You'll also pick up practical skills for cleaning and preparing data.

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.

The Architecture of Convnets and Data Processing.

Advantages of Convolutional Neural Networks - 2020-07-27
Essay from the year 2020 in the subject Mathematics - Miscellaneous, grade: 1,3, University of Ulm, language: English, abstract: In the past two decades in particular, artificial neural networks have led to new approaches and processes in machine learning in many areas. They have replaced many existing processes. In some areas, they even exceed human performance. Impressive progress has been made in the area of image recognition and classification. Above all, this includes the introduction of convolutional neural networks (ConvNets). They belong to the class of neural networks. The first ConvNet was developed by LeCun et al. in 1989. ConvNets were especially developed to enhance image processing. Therefore, they provide a unique architecture. Due to their structure and functionality, ConvNets are particularly well suited within this field of application compared to other methods.

Deep Learning for Robot Perception and Cognition - Alexandros Iosifidis 2022-02-25

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

Convolutional Neural Networks in Python - Anthony Williams 2017-07-25

Convolutional Neural Networks in Python (2nd Edition)

Deep learning has been a great part of various scientific fields and since this is my third book regarding this topic, you already know the great significance of deep learning in comparison to traditional methods. At this point, you are also familiar with types of neural networks and their wide range of applications including image and speech recognition, natural language processing, video game development and other. On the other hand, this book is all about convolutional neural networks and how to use these neural networks in various tasks of automatic image and speech recognition in Python. You will also get a better insight into the architecture of convolutional layers as we are going deeper into this subject. Deep learning is pretty complex subject, but since you already have a fundamental knowledge of this topic, getting to know convolutional neural networks better is next logical step. What you will learn in Convolutional Neural Networks in Python: Architecture of convolutional neural networks Solving computer vision tasks using convolutional neural networks Python and computer vision Automatic image and speech recognition Theano and TenorFlow image recognition How to use MNIST vision dataset What are commonly used convolutional filters Get this book today and learn more about Convolutional Neural Networks in Python!! PS: Get the Paperback and get this Ebook for FREE!!

Make Your Own Neural Network: An In-Depth Visual

Introduction for Beginners - Michael Taylor 2017-10-04
A step-by-step visual journey through the mathematics of neural networks, and making your own using Python and Tensorflow. What you will gain from this book: * A deep understanding of how a Neural Network works. * How to build a Neural Network from scratch using Python. Who this book is for: * Beginners who want to fully understand how networks work, and learn to build two step-by-step examples in Python. * Programmers who need an easy to read, but solid refresher, on the math of neural networks. What's Inside - 'Make Your Own Neural Network: An Indepth Visual Introduction For Beginners' What Is a Neural Network? Neural networks have made a gigantic comeback in the last few decades and you likely make use of them everyday without realizing it, but what exactly is a neural network? What is it used for and how does it fit within the broader arena of machine learning? we gently explore these topics so that we can be prepared to dive deep further on. To start, we'll begin with a high-level overview of machine learning and then drill down into the specifics of a neural network. The Math of Neural Networks On a high level, a network learns just like we do, through trial and error. This is true regardless if the network is supervised, unsupervised, or semi-supervised. Once we dig a bit deeper though, we discover that a handful of mathematical functions play a major role in the trial and error process. It also becomes clear that a grasp of the underlying mathematics helps clarify how a network learns. * Forward Propagation * Calculating The Total Error * Calculating The Gradients * Updating The Weights Make Your Own Artificial Neural Network: Hands on Example You will learn to build a simple neural network using all the concepts and functions we learned in the previous few chapters. Our example will be basic but hopefully very intuitive. Many examples available online are either hopelessly abstract or make use of the same data sets, which can be repetitive. Our goal is to be crystal clear and engaging, but with a touch of fun and uniqueness. This section contains the following eight chapters. Building Neural Networks in Python There are many ways to build a neural network and lots of tools to get the job done. This is fantastic, but it can also be overwhelming when you start, because there are so many tools to choose from. We are going to take a look at what tools are needed and help you nail down the essentials. To build a neural network Tensorflow and

Neural Networks There is no single way to build a feedforward neural network with Python, and that is especially true if you throw Tensorflow into the mix. However, there is a general framework that exists that can be divided into five steps and grouped into two parts. We are going to briefly explore these five steps so that we are prepared to use them to build a network later on. Ready? Let's begin. Neural Network:

Distinguish Handwriting We are going to dig deep with Tensorflow and build a neural network that can distinguish between handwritten numbers. We'll use the same 5 steps we covered in the high-level overview, and we are going to take time exploring each line of code.
Neural Network: Classify Images 10 minutes. That's all it takes to build an image classifier thanks to Google! We will provide a high-level overview of how to classify images using a convolutional neural network (CNN) and Google's Inception V3 model. Once finished, you will be able to tweak this code to classify any type of image sets! Cats, bats, super heroes - the sky's the limit.

State of the Art in Neural Networks and Their Applications - Ayman S. El-Baz 2021-07-21

State of the Art in Neural Networks and Their Applications presents the latest advances in artificial neural networks and their applications across a wide range of clinical diagnoses. Advances in the role of machine learning, artificial intelligence, deep learning, cognitive image processing and suitable data analytics useful for clinical diagnosis and research applications are covered, including relevant case studies. The application of Neural Network, Artificial Intelligence, and Machine Learning methods in biomedical image analysis have resulted in the development of computer-aided diagnostic (CAD) systems that aim towards the automatic early detection of several severe diseases. **State of the Art in Neural Networks and Their Applications** is presented in two volumes. Volume 1 covers the state-of-the-art deep learning approaches for the detection of renal, retinal, breast, skin, and dental abnormalities and more. Includes applications of neural networks, AI, machine learning, and deep learning techniques to a variety of imaging technologies Provides in-depth technical coverage of computer-aided diagnosis (CAD), with coverage of computer-aided classification, Unified Deep Learning Frameworks, mammography, fundus imaging, optical coherence tomography, cryo-electron tomography, 3D MRI, CT, and more. Covers deep learning for several medical conditions including renal, retinal, breast, skin, and dental abnormalities, Medical Image Analysis, as well as detection, segmentation, and classification via AI.

Deep Learning With Python - Jason Brownlee 2016-05-13
Deep learning is the most interesting and powerful machine learning technique right now. Top deep learning libraries are available on the Python ecosystem like Theano and TensorFlow. Tap into their power in a few lines of code using Keras, the best-of-breed applied deep learning library. In this Ebook, learn exactly how to get started and apply deep learning to your own machine learning projects.

Tree-Based Convolutional Neural Networks - Lili Mou 2018-10-01

This book proposes a novel neural architecture, tree-based convolutional neural networks (TBCNNs), for processing tree-structured data. TBCNNs are related to existing convolutional neural networks (CNNs) and recursive neural networks (RNNs), but they combine the merits of both: thanks to their short propagation path, they are as efficient in learning as CNNs; yet they are also as structure-sensitive as RNNs. In this book, readers will also find a comprehensive literature review of related work, detailed descriptions of TBCNNs and their variants, and experiments applied to program analysis and natural language processing tasks. It is also an enjoyable read for all those with a general

interest in deep learning.

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

Introduction to Deep Learning - Eugene Charniak 2019-01-29

A project-based guide to the basics of deep learning. This concise, project-driven guide to deep learning takes readers through a series of program-writing tasks that introduce them to the use of deep learning in such areas of artificial intelligence as computer vision, natural-language processing, and reinforcement learning. The author, a longtime artificial intelligence researcher specializing in natural-language processing, covers feed-forward neural nets, convolutional neural nets, word embeddings, recurrent neural nets, sequence-to-sequence learning, deep reinforcement learning, unsupervised models, and other fundamental concepts and techniques. Students and practitioners learn the basics of deep learning by working through programs in Tensorflow, an open-source machine learning framework. "I find I learn computer science material best by sitting down and writing programs," the author writes, and the book reflects this approach. Each chapter includes a programming project, exercises, and references for further reading. An early chapter is devoted to Tensorflow and its interface with Python, the widely used programming language. Familiarity with linear algebra, multivariate calculus, and probability and statistics is required, as is a rudimentary knowledge of programming in Python. The book can be used in both undergraduate and graduate courses; practitioners will find it an essential reference.
Machine Learning and Deep Learning in Real-Time Applications - Mahrishi, Mehul 2020-04-24
Artificial intelligence and its various components are rapidly engulfing almost every professional industry. Specific features of AI that have proven to be vital solutions to numerous real-world issues are machine learning and deep learning. These intelligent agents unlock higher levels of performance and efficiency, creating a wide span of industrial applications. However, there is a lack of research on the specific uses of machine/deep learning in the professional realm. **Machine Learning and Deep Learning in Real-Time Applications** provides emerging research exploring the theoretical and practical aspects of machine learning and deep learning and their implementations as well as their ability to solve real-world problems within several professional disciplines including healthcare, business, and computer science. Featuring coverage on a

broad range of topics such as image processing, medical improvements, and smart grids, this book is ideally designed for researchers, academicians, scientists, industry experts, scholars, IT professionals, engineers, and students seeking current research on the multifaceted uses and implementations of machine learning and deep learning across the globe.

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

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.

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 ...

Neural Networks and Statistical Learning - Ke-Lin Du 2013-12-09

Providing a broad but in-depth introduction to neural network and machine learning in a statistical framework, this book provides a single, comprehensive resource for study and further research. All the major popular neural network models and statistical learning approaches are covered with examples and exercises in every chapter to develop a practical working understanding of the content. Each of the twenty-five chapters includes state-of-the-art descriptions and important research results on the respective topics. The broad coverage includes the multilayer perceptron, the Hopfield network, associative memory models, clustering models and algorithms, the radial basis function network, recurrent neural networks, principal component analysis, nonnegative matrix factorization, independent component analysis, discriminant analysis, support vector machines, kernel methods, reinforcement learning, probabilistic and Bayesian networks, data fusion and ensemble learning, fuzzy sets and logic, neurofuzzy models, hardware implementations, and some machine learning topics. Applications to biometric/bioinformatics and data mining are also included. Focusing on the prominent accomplishments and their practical aspects, academic and technical staff, graduate students and researchers will find that this provides a solid foundation and encompassing reference for the fields of neural networks, pattern recognition, signal processing, machine learning, computational intelligence, and data mining.

Introduction to Deep Learning - Sandro Skansi 2018-02-04

This textbook presents a concise, accessible and engaging first introduction to deep learning, offering a wide range of connectionist models which represent the current state-of-the-art. The text explores the most popular algorithms and architectures in a simple and intuitive style, explaining the mathematical derivations in a step-by-step manner. The content coverage includes convolutional networks, LSTMs, Word2vec, RBMs, DBNs, neural Turing machines, memory networks and autoencoders. Numerous examples in working Python code are provided throughout the book, and the code is also supplied separately at an accompanying website. Topics and features: introduces the fundamentals of machine learning, and the mathematical and computational

prerequisites for deep learning; discusses feed-forward neural networks, and explores the modifications to these which can be applied to any neural network; examines convolutional neural networks, and the recurrent connections to a feed-forward neural network; describes the notion of distributed representations, the concept of the autoencoder, and the ideas behind language processing with deep learning; presents a brief history of artificial intelligence and neural networks, and reviews interesting open research problems in deep learning and connectionism. This clearly written and lively primer on deep learning is essential reading for graduate and advanced undergraduate students of computer science, cognitive science and mathematics, as well as fields such as linguistics, logic, philosophy, and psychology.

Deep Learning for Coders with fastai and PyTorch -

Jeremy Howard 2020-06-29

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

Deep Learning for Computer Vision - Jason Brownlee
2019-04-04

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

Introduction to Graph Neural Networks - Zhiyuan Zhiyuan Liu 2022-05-31

Graphs are useful data structures in complex real-life applications such as modeling physical systems, learning molecular fingerprints, controlling traffic networks, and recommending friends in social networks. However, these tasks require dealing with non-Euclidean graph data that contains rich relational information between elements and cannot be well handled by traditional deep learning models (e.g., convolutional neural networks (CNNs) or recurrent neural networks (RNNs)). Nodes in graphs usually contain useful feature information that cannot be well addressed in most unsupervised representation learning methods (e.g., network embedding methods). Graph neural networks (GNNs) are proposed to combine the feature information and the graph structure to learn better representations on graphs via feature propagation and aggregation. Due to its convincing performance and high interpretability, GNN has recently become a widely applied graph analysis tool. This book provides a comprehensive introduction to the basic concepts, models, and applications of graph neural networks. It starts with the introduction of the vanilla GNN model. Then several variants of the vanilla model are introduced such as graph convolutional networks, graph recurrent networks, graph attention networks, graph residual networks, and several general frameworks. Variants for different graph types and advanced training methods are also included. As for the applications of GNNs, the book categorizes them into structural, non-structural, and other scenarios, and then it introduces

several typical models on solving these tasks. Finally, the closing chapters provide GNN open resources and the outlook of several future directions.

Hands-On Convolutional Neural Networks with TensorFlow - Iffat Zafar 2018-08-28

Learn how to apply TensorFlow to a wide range of deep learning and Machine Learning problems with this practical guide on training CNNs for image classification, image recognition, object detection and many computer vision challenges. Key Features Learn the fundamentals of Convolutional Neural Networks Harness Python and Tensorflow to train CNNs Build scalable deep learning models that can process millions of items Book Description Convolutional Neural Networks (CNN) are one of the most popular architectures used in computer vision apps. This book is an introduction to CNNs through solving real-world problems in deep learning while teaching you their implementation in popular Python library - TensorFlow. By the end of the book, you will be training CNNs in no time! We start with an overview of popular machine learning and deep learning models, and then get you set up with a TensorFlow development environment. This environment is the basis for implementing and training deep learning models in later chapters. Then, you will use Convolutional Neural Networks to work on problems such as image classification, object detection, and semantic segmentation. After that, you will use transfer learning to see how these models can solve other deep learning problems. You will also get a taste of implementing generative models such as autoencoders and generative adversarial networks. Later on, you will see useful tips on machine learning best practices and troubleshooting. Finally, you will learn how to apply your models on large datasets of millions of images. What you will learn Train machine learning models with TensorFlow Create systems that can evolve and scale during their life cycle Use CNNs in image recognition and classification Use TensorFlow for building deep learning models Train popular deep learning models Fine-tune a neural network to improve the quality of results with transfer learning Build TensorFlow models that can scale to large datasets and systems Who this book is for This book is for Software Engineers, Data Scientists, or Machine Learning practitioners who want to use CNNs for solving real-world problems. Knowledge of basic machine learning concepts, linear algebra and Python will help. **Machine Learning with Neural Networks** - Bernhard Mehlig 2021-08-31

This modern and self-contained book offers a clear and accessible introduction to the important topic of machine learning with neural networks. In addition to describing the mathematical principles of the topic, and its historical evolution, strong connections are drawn with underlying methods from statistical physics and current applications within science and engineering. Closely based around a well-established undergraduate course, this pedagogical text provides a solid understanding of the key aspects of modern machine learning with artificial neural networks, for students in physics, mathematics, and engineering. Numerous exercises expand and reinforce key concepts within the book and allow students to hone their programming skills. Frequent references to current research develop a detailed perspective on the state-of-the-art in machine learning research.

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 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.

Images as Data for Social Science Research - Nora Webb Williams 2020-08-31

Images play a crucial role in shaping and reflecting political life. Digitization has vastly increased the presence of such images in daily life, creating valuable new research opportunities for social scientists. We show how recent innovations in computer vision methods can substantially lower the costs of using images as data. We introduce readers to the deep learning algorithms commonly used for object recognition, facial recognition, and visual sentiment analysis. We then provide guidance and specific instructions for scholars interested in using these methods in their own research.

Deep Learning for Time Series Forecasting - Jason Brownlee 2018-08-30

Deep learning methods offer a lot of promise for time series forecasting, such as the automatic learning of temporal dependence and the automatic handling of temporal structures like trends and seasonality. With clear explanations, standard Python libraries, and step-by-step tutorial lessons you'll discover how to develop deep learning models for your own time series forecasting projects.

Practical Deep Learning - Ronald T. Kneusel 2021-02-23 Practical Deep Learning teaches total beginners how to build the datasets and models needed to train neural networks for your own DL projects. If you've been curious about machine learning but didn't know where to start, this is the book you've been waiting for. Focusing on the subfield of machine learning known as deep learning, it explains core concepts and gives you the foundation you need to start building your own models. Rather than simply outlining recipes for using existing toolkits, Practical Deep Learning teaches you the why of deep learning and will inspire you to explore further. All you need is basic familiarity with computer programming and high school math—the book will cover the rest. After an introduction to Python, you'll move through key topics like how to build a good training dataset, work with the scikit-learn and Keras libraries, and evaluate your models' performance. You'll also learn: How to use classic machine learning models like k-Nearest Neighbors, Random Forests, and Support Vector Machines How neural networks work and how they're trained How to use convolutional neural networks How to develop a successful deep learning model from scratch You'll conduct experiments along the way, building to a final case study that incorporates everything you've learned. The perfect introduction to this dynamic, ever-

expanding field, Practical Deep Learning will give you the skills and confidence to dive into your own machine learning projects.

Long Short-Term Memory Networks With Python - Jason Brownlee 2017-07-20

The Long Short-Term Memory network, or LSTM for short, is a type of recurrent neural network that achieves state-of-the-art results on challenging prediction problems. In this laser-focused Ebook, finally cut through the math, research papers and patchwork descriptions about LSTMs. Using clear explanations, standard Python libraries and step-by-step tutorial lessons you will discover what LSTMs are, and how to develop a suite of LSTM models to get the most out of the method on your sequence prediction problems.

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.

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.

Neural Networks and Convolutional Neural Networks Essential Training - Jonathan Fernandes 2018

Images as Data for Social Science Research - Nora Webb Williams 2020-08-31

Images play a crucial role in shaping and reflecting political life. Digitization has vastly increased the presence of such images in daily life, creating valuable

new research opportunities for social scientists. We show how recent innovations in computer vision methods can substantially lower the costs of using images as data. We introduce readers to the deep learning algorithms commonly used for object recognition, facial recognition, and visual sentiment analysis. We then provide guidance and specific instructions for scholars interested in using these methods in their own research.

Deep Learning - Ian Goodfellow 2016-11-10

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

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 Python - Nikhil Ketkar 2017-04-18
Discover the practical aspects of implementing deep-learning solutions using the rich Python ecosystem. This book bridges the gap between the academic state-of-the-art and the industry state-of-the-practice by introducing you to deep learning frameworks such as Keras, Theano, and Caffe. The practicalities of these frameworks is often acquired by practitioners by reading

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

Neural Networks with R - Giuseppe Ciaburro 2017-09-27

Uncover the power of artificial neural networks by implementing them through R code. About This Book Develop a strong background in neural networks with R, to implement them in your applications Build smart systems using the power of deep learning Real-world case studies to illustrate the power of neural network models Who This Book Is For This book is intended for anyone who has a statistical background with knowledge in R and wants to work with neural networks to get better results from complex data. If you are interested in artificial intelligence and deep learning and you want to level up, then this book is what you need! What You Will Learn Set up R packages for neural networks and deep learning Understand the core concepts of artificial neural networks Understand neurons, perceptrons, bias, weights, and activation functions Implement supervised and unsupervised machine learning in R for neural networks Predict and classify data automatically using neural networks Evaluate and fine-tune the models you build. In Detail Neural networks are one of the most fascinating machine learning models for solving complex computational problems efficiently. Neural networks are used to solve wide range of problems in different areas of AI and machine learning. This book explains the niche aspects of neural networking and provides you with foundation to get started with advanced topics. The book begins with neural network design using the neural net package, then you'll build a solid foundation knowledge of how a neural network learns from data, and the principles behind it. This book covers various types of neural network including recurrent neural networks and convoluted neural networks. You will not only learn how to train neural networks, but will also explore generalization of these networks. Later we will delve into combining different neural network models and work with the real-world use cases. By the end of this book, you will learn to implement neural network models in your applications with the help of practical examples in the book. Style and approach A step-by-step guide filled with real-world practical examples.

Deep Learning on Graphs - Yao Ma 2021-09-23

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