

An Introduction To Convolutional Neural Networks

[Java Deep Learning Cookbook](#)
[State of the Art in Neural Networks and Their Applications](#)
[Machine Learning and Deep Learning in Real-Time Applications](#)
[Introduction to Deep Learning](#)
[Guide to Convolutional Neural Networks](#)
[Proceedings of ICTIS 2020, Volume 1](#)
[Tree-Based Convolutional Neural Networks](#)
[Introduction to Convolutional Neural Networks](#)
[Deep Learning and Convolutional Neural Networks for Medical Image Computing](#)
[The Architecture of Convnets and Data Processing. Advantages of Convolutional Neural Networks Volume 1](#)
[A Guide to Convolutional Neural Networks for Computer Vision](#)
[Images as Data for Social Science Research](#)
[A Practical Application to Traffic-Sign Detection and Classification](#)
[Artificial Intelligence and Global Society](#)
[Introduction to Deep Learning for Healthcare](#)
[Implement advanced deep learning models using Python](#)
[Precision Medicine, High Performance and Large-Scale Datasets](#)
[A Practical Guide](#)
[A Concise Guide](#)
[Deep Learning for Computer Vision](#)
[Efficient Processing of Deep Neural Networks](#)
[Proceedings of the 1st International Conference on Electronics, Biomedical Engineering, and Health Informatics](#)
[Machine Learning with Neural Networks](#)
[Train neural networks for classification, NLP, and reinforcement learning using Deeplearning4j](#)
[A Textbook](#)
[Deep Learning on Graphs](#)
[Deep Learning with Python](#)
[From Logical Calculus to Artificial Intelligence](#)
[A Greater Foundation for Machine Learning Engineering](#)
[Practical Convolutional Neural Networks](#)
[An Introduction for Scientists and Engineers](#)
[A Guide to Building Deep Learning Systems](#)
[The Hallmarks of the Great Beyond in Pytorch, R, Tensorflow, and Python](#)
[An Introduction to Convolutional Neural Nets for Image Classification](#)
[Convolutional Neural Networks and Object Detection](#)
[Convolutional Neural Networks in Visual Computing](#)
[Convolutional Neural Networks in Python](#)
[Using Convolutional Neural Networks and TensorFlow to Solve Computer Vision Tasks](#)
[Advanced Applied Deep Learning](#)

An Introduction To Convolutional Neural Networks

Downloaded from [archive.imba.com](#) by guest

BENTON MADELINE

Java Deep Learning Cookbook Packt Publishing Ltd

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.

State of the Art in Neural Networks and Their Applications CRC Press

Discover the potential applications, challenges, and opportunities of deep learning from a business perspective with technical examples. These applications include image recognition, segmentation and annotation, video processing and annotation, voice recognition, intelligent personal assistants, automated translation, and autonomous vehicles. An Introduction to Deep Learning Business Applications for Developers covers some common DL algorithms such as content-based recommendation algorithms and natural language processing. You'll explore examples, such as video prediction with fully convolutional neural networks (FCNN) and residual neural networks (ResNets). You will also see applications of DL for controlling robotics, exploring the DeepQ learning algorithm with Monte Carlo Tree search (used to beat humans in the game of Go), and modeling for financial risk assessment. There will also be mention of the powerful set of algorithms called Generative Adversarial Neural networks (GANs) that can be applied for image colorization, image completion, and style transfer. After reading this book you will have an overview of the exciting field of deep neural networks and an understanding of most of the major applications of deep learning. The book contains some coding examples, tricks, and insights on how to train deep learning models using the Keras framework. What You Will Learn Find out about deep learning and why it is so powerful Work with the major algorithms available to train deep learning models See the major breakthroughs in terms of applications of deep learning Run simple examples with a selection of deep learning libraries Discover the areas of impact of deep learning in business Who This Book Is For Data scientists, entrepreneurs, and business developers.

Machine Learning and Deep Learning in Real-Time Applications Academic Press

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.

Introduction to Deep Learning "O'Reilly Media, Inc."

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 ...

Guide to Convolutional Neural Networks Academic Press

Introduction to Deep Learning and Neural Networks with Python™: A Practical Guide is an intensive step-by-step guide for neuroscientists to fully understand, practice, and build neural networks. Providing math and Python™ 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 Python™ 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 Python™ functions and features for neuroscientists Uses a careful tutorial approach to describe implementation of neural networks in Python™ Features math and code examples (via companion website) with helpful instructions for easy implementation

Proceedings of ICTIS 2020, Volume 1 Cambridge University Press

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.

Tree-Based Convolutional Neural Networks Packt Publishing Ltd

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.

Introduction to Convolutional Neural Networks Apress

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.

Deep Learning and Convolutional Neural Networks for Medical Image Computing MIT Press

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.

The Architecture of Convnets and Data Processing. Advantages of Convolutional Neural Networks CRC Press

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.

Volume 1 Apress

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

A Guide to Convolutional Neural Networks for Computer Vision Morgan & Claypool Publishers

It's time to dispel the myth that machine learning is difficult. Grokking Machine Learning teaches you how to apply ML to your projects using only standard Python code and high school-level math. No specialist knowledge is required to tackle the hands-on exercises using readily available machine learning tools! In Grokking Machine Learning, expert machine learning engineer Luis Serrano introduces the most valuable ML techniques and teaches you how to make them work for you. Practical examples illustrate each new concept to ensure you're grokking as you go. You'll build models for spam detection, language analysis, and image recognition as you lock in each carefully-selected skill. Packed with easy-to-follow Python-based exercises and mini-projects, this book sets you on the path to becoming a machine learning expert. Key Features · Different types of machine learning, including supervised and unsupervised learning · Algorithms for simplifying, classifying, and splitting data · Machine learning packages and tools · Hands-on exercises with fully-explained Python code samples For readers with intermediate programming knowledge in Python or a similar language. About the technology Machine learning is a collection of mathematically-based techniques and algorithms that enable computers to identify patterns and generate predictions from data. This

revolutionary data analysis approach is behind everything from recommendation systems to self-driving cars, and is transforming industries from finance to art.

Images as Data for Social Science Research Apress

Develop and optimize deep learning models with advanced architectures. This book teaches you the intricate details and subtleties of the algorithms that are at the core of convolutional neural networks. In Advanced Applied Deep Learning, you will study advanced topics on CNN and object detection using Keras and TensorFlow. Along the way, you will look at the fundamental operations in CNN, such as convolution and pooling, and then look at more advanced architectures such as inception networks, resnets, and many more. While the book discusses theoretical topics, you will discover how to work efficiently with Keras with many tricks and tips, including how to customize logging in Keras with custom callback classes, what is eager execution, and how to use it in your models. Finally, you will study how object detection works, and build a complete implementation of the YOLO (you only look once) algorithm in Keras and TensorFlow. By the end of the book you will have implemented various models in Keras and learned many advanced tricks that will bring your skills to the next level. What You Will Learn See how convolutional neural networks and object detection work Save weights and models on disk Pause training and restart it at a later stage Use hardware acceleration (GPUs) in your code Work with the Dataset TensorFlow abstraction and use pre-trained models and transfer learning Remove and add layers to pre-trained networks to adapt them to your specific project Apply pre-trained models such as Alexnet and VGG16 to new datasets Who This Book Is For Scientists and researchers with intermediate-to-advanced Python and machine learning know-how. Additionally, intermediate knowledge of Keras and TensorFlow is expected.

A Practical Application to Traffic-Sign Detection and Classification Springer

In this video course, you will learn the basic principles of neural networks that are used to build models. You'll start by seeing machine learning, neurons, activations, activation functions, weights, and how everything works under the hood. Next, you'll cover the basics of the learning loop including how backpropagation and gradient descent work. Further, you will learn about convolutions, how they are inspired by the animal visual cortex, and how we use them in neural networks. One of the focuses of the course is image classification and detecting common objects in images. This has many uses in your day-to-day projects. We will be using the PyTorch open-source neural network library here. The course will also cover current state-of-the-art neural network models and show how to use them even on smaller hardware. The video concludes by showing some common tricks with hyperparameter settings and regularization techniques, and how to use neural networks in production environments. What You Will Learn Discover the basics of neural networks and how they function Work with convolutional neural networks Use CNNs in your day-to-day work for image classification and other tasks Who This Video Is For Data scientists and machine learning and deep learning engineers.

Artificial Intelligence and Global Society Springer Nature

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 for Healthcare Morgan & Claypool Publishers

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.

Implement advanced deep learning models using Python BPB Publications

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.

Precision Medicine, High Performance and Large-Scale Datasets Images as Data for Social Science Research
 An Introduction to Convolutional Neural Nets for Image Classification
 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.

A Practical Guide Springer

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.

A Concise Guide Lulu.com

"This course provides an introduction to using deep learning models to solve computer vision tasks in TensorFlow and is focused on the work horse of deep learning image models: the convolutional neural network. Expert Lucas Adams teaches you how to get these models up and running fast, especially in domains with limited computing resources or training data, and shows you how to modify the architecture of a neural network to make the model specialized to different tasks. Learners should be familiar with basic deep learning concepts like the multilayer perceptron, linear algebra, Jupyter notebooks, and the basics of building and running TensorFlow programs."--Resource description page.

Related with An Introduction To Convolutional Neural Networks:

- Economics Of Ancient Mesopotamia : [click here](#)