

Stanford University Tensorflow For Deep Learning Research

Leverage the power of Tensorflow to design deep learning systems for a variety of real-world scenarios Key Features Build efficient deep learning pipelines using the popular Tensorflow framework Train neural networks such as ConvNets, generative models, and LSTMs Includes projects related to Computer Vision, stock prediction, chatbots and more Book Description TensorFlow is one of the most popular frameworks used for machine learning and, more recently, deep learning. It provides a fast and efficient framework for training different kinds of deep learning models, with very high accuracy. This book is your guide to master deep learning with TensorFlow with the help of 10 real-world projects. TensorFlow Deep Learning Projects starts with setting up the right TensorFlow environment for deep learning. Learns to train different types of deep learning models using TensorFlow, including Convolutional Neural Networks, Recurrent Neural Networks, LSTMs, and Generative Adversarial Networks. While doing so, you will build end-to-end deep learning solutions to tackle different real-world problems in image processing, recommendation systems, stock prediction, and building chatbots, to name a few. You will also develop systems that perform machine translation, and use reinforcement learning techniques to play games. By the end of this book, you will have mastered all the concepts of deep learning and their implementation with TensorFlow, and understand the entire system from hardware to software. The purpose of this book is to provide a solid understanding of (1) the design, training, and applications of DL algorithms in industry; (2) the compiler techniques to map deep learning code to hardware targets; and (3) the critical hardware features that accelerate DL systems. This book aims to facilitate co-innovation for the advancement of DL systems. It is written for engineers working in one or more of these areas who seek to understand the entire system from hardware to software in order to better collaborate with engineers working in other parts of the system stack. The book details advancements and adoption of DL models in industry, explains the training and deployment process, describes the essential hardware architectural features needed for today's and future models, and details advances in DL compilers to efficiently execute algorithms across various hardware targets. Unique in this book is the holistic exposition of the entire DL system stack, the emphasis on commercial applications, and the practical techniques to design models and accelerate their performance. The author is fortunate to work with hardware, software, data scientist, and research teams across many high-technology companies with hyperscale data centers. These companies employ many of the examples and methods provided throughout the book. Explore machine learning concepts using the latest numerical computing library — TensorFlow — with the help of this comprehensive cookbook About This Book Your quick guide to implementing TensorFlow in your day-to-day machine learning activities Learn advanced techniques that bring more accuracy and speed to machine learning Upgrade your knowledge to the second generation of machine learning with this guide on TensorFlow Who This Book Is For This book is ideal for data scientists who are familiar with C++ or Python and perform machine learning utilities on a day-to-day basis. Intermediate and advanced machine learning implementers who need a quick guide they can easily navigate will find it useful. What You Will Learn Become familiar with the basics of the TensorFlow machine learning library Get to know Linear Regression techniques with TensorFlow Learn SVMs with hands-on recipes Implement neural networks and improve predictions Apply NLP and sentiment analysis to your data Master CNN and RNN through practical recipes Take TensorFlow into production In Detail TensorFlow is an open source software library for Machine Intelligence. The independent recipes in this book will teach you how to use TensorFlow for complex data computations and will let you dig deeper and gain more insights into your data than ever before. You'll work through recipes on training models, model evaluation, sentiment analysis, regression analysis, and more. TensorFlow is a deep learning model with TensorFlow confidently. What you will learn Set up the TensorFlow environment for deep learning Construct your own ConvNets for effective image processing Use LSTM for image caption generation Forecast stock prediction accurately with an LSTM architecture Learn what semantic matching is by detecting duplicate photos questions Set up an AWS instance with TensorFlow to train GANs Train and set up a chatbot to understand and interpret human input Build an AI capable of playing a video game by itself — and win it! Who this book is for This book is for data scientists, machine learning developers as well as deep learning practitioners, who want to build interesting deep learning projects that leverage the power of TensorFlow. Some understanding of machine learning and deep learning, and familiarity with the TensorFlow framework is all you need to get started with this book. This book describes deep learning systems: the algorithms, compilers, and processor components to efficiently train and deploy deep learning models for commercial applications. The exponential growth in computational power is slowing at a time when the amount of compute consumed by state-of-the-art deep learning (DL) workloads is rapidly growing. Model size, serving latency, and power constraints are a significant challenge in the deployment of DL models for many applications. Therefore, it is imperative to co-design algorithms, compilers, and hardware to accelerate advances in this field with holistic system-level and algorithm solutions that improve performance, power, and efficiency. Advancing DL systems generally involves three types of engineers: (1) data scientists that utilize and develop DL algorithms in partnership with domain experts, such as medical, economic, or climate scientists; (2) hardware designers that develop specialized hardware to accelerate the components in the DL models; and (3) performance and compiler engineers who optimize software to run more efficiently on a given hardware. Hardware engineers should be aware of the characteristics and components of production and academic models likely to be adopted by industry to guide design decisions impacting future hardware. Data scientists should be aware of deployment power constraints when designing models. Performance engineers should support optimizations across diverse devices, libraries, and hardware targets. The purpose of this book is to provide a solid understanding of (1) the design, training, and applications of DL algorithms in industry; (2) the compiler techniques to map deep learning code to hardware targets; and (3) the critical hardware features that accelerate DL systems. This book aims to facilitate co-innovation for the advancement of DL systems. It is written for engineers working in one or more of these areas who seek to understand the entire system from hardware to software in order to better collaborate with engineers working in other parts of the system stack. The book details advancements and adoption of DL models in industry, explains the training and deployment process, describes the essential hardware architectural features needed for today's and future models, and details advances in DL compilers to efficiently execute algorithms across various hardware targets. Unique in this book is the holistic exposition of the entire DL system stack, the emphasis on commercial applications, and the practical techniques to design models and accelerate their performance. The author is fortunate to work with hardware, software, data scientist, and research teams across many high-technology companies with hyperscale data centers. These companies employ many of the examples and methods provided throughout the book.

Explore machine learning concepts using the latest numerical computing library — TensorFlow — with the help of this comprehensive cookbook About This Book Your quick guide to implementing TensorFlow in your day-to-day machine learning activities Learn advanced techniques that bring more accuracy and speed to machine learning Upgrade your knowledge to the second generation of machine learning with this guide on TensorFlow Who This Book Is For This book is ideal for data scientists who are familiar with C++ or Python and perform machine learning utilities on a day-to-day basis. Intermediate and advanced machine learning implementers who need a quick guide they can easily navigate will find it useful. What You Will Learn Become familiar with the basics of the TensorFlow machine learning library Get to know Linear Regression techniques with TensorFlow Learn SVMs with hands-on recipes Implement neural networks and improve predictions Apply NLP and sentiment analysis to your data Master CNN and RNN through practical recipes Take TensorFlow into production In Detail TensorFlow is an open source software library for Machine Intelligence. The independent recipes in this book will teach you how to use TensorFlow for complex data computations and will let you dig deeper and gain more insights into your data than ever before. You'll work through recipes on training models, model evaluation, sentiment analysis, regression analysis, and more. TensorFlow is a deep learning model with TensorFlow confidently. What you will learn Set up the TensorFlow environment for deep learning Construct your own ConvNets for effective image processing Use LSTM for image caption generation Forecast stock prediction accurately with an LSTM architecture Learn what semantic matching is by detecting duplicate photos questions Set up an AWS instance with TensorFlow to train GANs Train and set up a chatbot to understand and interpret human input Build an AI capable of playing a video game by itself — and win it! Who this book is for This book is for data scientists, machine learning developers as well as deep learning practitioners, who want to build interesting deep learning projects that leverage the power of TensorFlow. Some understanding of machine learning and deep learning, and familiarity with the TensorFlow framework is all you need to get started with this book. This book describes deep learning systems: the algorithms, compilers, and processor components to efficiently train and deploy deep learning models for commercial applications. The exponential growth in computational power is slowing at a time when the amount of compute consumed by state-of-the-art deep learning (DL) workloads is rapidly growing. Model size, serving latency, and power constraints are a significant challenge in the deployment of DL models for many applications. Therefore, it is imperative to co-design algorithms, compilers, and hardware to accelerate advances in this field with holistic system-level and algorithm solutions that improve performance, power, and efficiency. Advancing DL systems generally involves three types of engineers: (1) data scientists that utilize and develop DL algorithms in partnership with domain experts, such as medical, economic, or climate scientists; (2) hardware designers that develop specialized hardware to accelerate the components in the DL models; and (3) performance and compiler engineers who optimize software to run more efficiently on a given hardware. Hardware engineers should be aware of the characteristics and components of production and academic models likely to be adopted by industry to guide design decisions impacting future hardware. Data scientists should be aware of deployment power constraints when designing models. Performance engineers should support optimizations across diverse devices, libraries, and hardware targets. The purpose of this book is to provide a solid understanding of (1) the design, training, and applications of DL algorithms in industry; (2) the compiler techniques to map deep learning code to hardware targets; and (3) the critical hardware features that accelerate DL systems. This book aims to facilitate co-innovation for the advancement of DL systems. It is written for engineers working in one or more of these areas who seek to understand the entire system from hardware to software in order to better collaborate with engineers working in other parts of the system stack. The book details advancements and adoption of DL models in industry, explains the training and deployment process, describes the essential hardware architectural features needed for today's and future models, and details advances in DL compilers to efficiently execute algorithms across various hardware targets. Unique in this book is the holistic exposition of the entire DL system stack, the emphasis on commercial applications, and the practical techniques to design models and accelerate their performance. The author is fortunate to work with hardware, software, data scientist, and research teams across many high-technology companies with hyperscale data centers. These companies employ many of the examples and methods provided throughout the book.

High-Performance Computing for Big Data: Methodologies and Applications explores emerging high-performance architectures for data-intensive applications, novel efficient analytical strategies to boost data processing, and cutting-edge applications in diverse fields, such as machine learning, life science, neural networks, and neuromorphic engineering. The book is organized into two main sections. The first section covers Big Data architectures, including cloud computing systems, and heterogeneous accelerators. It also covers emerging 3D IC design principles for memory architectures and devices. The second section of the book illustrates emerging and practical applications of Big Data across several domains, including bioinformatics, deep learning, and neuromorphic engineering. Features Covers a wide range of Big Data architectures, including distributed systems like Hadoop/Spark Includes accelerator-based approaches for big data applications such as GPU-based acceleration techniques, and hardware acceleration such as FPGA/CGRA/ASICs Presents emerging memory architectures and devices such as NVM, STT- RAM, 3D IC design principles Describes advanced algorithms for different big data application domains Illustrates novel analytics techniques for Big Data applications, scheduling, mapping, and partitioning methodologies Featuring contributions from leading experts, this book presents state-of-the-art research on the methodologies and applications of high-performance computing for big data applications. About the Editor Dr. Chao Wang is an Associate Professor in the School of Computer Science at the University of Science and Technology of China. He is the Associate Editor of ACM Transactions on Design Automations for Electronics Systems (TODAES), Applied Soft Computing, Microprocessors and Microsystems, IET Computers & Digital Techniques, and International Journal of Electronics. Dr. Chao Wang was the recipient of Youth Innovation Promotion Association, CAS, ACM China Rising Star Honorable Mention (2016), and best IP nomination of DATE 2015. He is now on the CCF Technical Committee on Computer Architecture, CCF Task Force on Formal Methods.

Reinforcement learning encompasses both a science of adaptive behavior of rational beings in uncertain environments and a computational methodology for finding optimal behaviors for challenging problems in control, optimization and adaptive behavior of intelligent agents. As a field, reinforcement learning has progressed tremendously in the past decade. The main goal of this book is to present an up-to-date series of survey articles on the main contemporary sub-fields of reinforcement learning. This includes surveys on partially observable environments, hierarchical task decompositions, relational knowledge representation and predictive state representations. Furthermore, topics such as transfer, evolutionary methods and continuous spaces in reinforcement learning are surveyed. In addition, several chapters review reinforcement learning methods in robotics, in games, and in computational neuroscience. In total seventeen different subfields are presented by mostly young experts in those areas, and together they truly represent a state-of-the-art of current reinforcement learning research. Marco Wiering works at the artificial intelligence department of the University of Groningen in the Netherlands. He has published extensively on various reinforcement learning topics. Martijn van Otterlo works in the cognitive artificial intelligence group at the Radboud University Nijmegen in The Netherlands. He has mainly focused on expressive knowledge representation in reinforcement learning settings.

Deploy deep learning solutions in production with ease using TensorFlow. You'll also develop the mathematical understanding and intuition required to invent new deep learning architectures and solutions on your own. Pro Deep Learning with TensorFlow provides practical, hands-on expertise so you can learn deep learning from scratch and deploy meaningful deep learning solutions. This book will allow you to get up to speed quickly using TensorFlow and to optimize different deep learning architectures. All of the practical aspects of deep learning that are relevant in any industry are emphasized in this book. You will be able to use the prototypes demonstrated to build new deep learning applications. The code presented in the book is available in the form of iPython notebooks and scripts which allow you to try out examples and extend them in interesting ways. You will be equipped with the mathematical foundation and scientific knowledge to pursue research in this field and give back to the community. What You'll Learn Understand full stack deep learning using TensorFlow and gain a solid mathematical foundation for deep learning Deploy complex deep learning solutions in production using TensorFlow Carry out research on deep learning and perform experiments using TensorFlow Who This Book Is For Data scientists and machine learning professionals, software developers, graduate students, and open source enthusiasts

Discover powerful ways to effectively solve real-world machine learning problems using key libraries including scikit-learn, TensorFlow, and PyTorch Key Features Learn and implement machine learning algorithms in a variety of real-life scenarios Cover a range of tasks catering to supervised, unsupervised and reinforcement learning techniques Find easy-to-follow code solutions for tackling common and not-so-common challenges Book Description This eagerly anticipated second edition of the popular Python Machine Learning Cookbook will enable you to adopt a fresh approach to dealing with real-world machine learning and deep learning tasks. With the help of over 100 recipes, you will learn to build powerful machine learning applications using modern libraries from the Python ecosystem. The book will also guide you on how to implement various machine learning algorithms for classification, clustering, and recommendation engines, using a recipe-based approach. With emphasis on practical solutions, dedicated sections in the book will help you to apply supervised and unsupervised learning techniques to real-world problems. Toward the concluding chapters, you will get to grips with recipes that teach you advanced techniques including reinforcement learning, deep neural networks, and automated machine learning. By the end of this book, you will be equipped with the skills you need to apply machine learning techniques and leverage the full capabilities of the Python ecosystem through real-world examples. What you will learn Use predictive modeling and apply it to a real-world problem using Epsilon-Shrink visualization techniques to interact with your data Learn how to build a recommendation engine Understand how to interact with text data and build models to analyze it Work with speech data and recognize spoken words using Hidden Markov Models Get well versed with reinforcement learning, automated ML, and transfer learning Work with image data and build systems for image recognition and biometric face recognition Use deep neural networks to build an optical character recognition system Who this book is for This book is for data scientists, machine learning developers, deep learning enthusiasts and Python programmers who want to solve real-world challenges using machine-learning techniques and algorithms. If you are facing challenges at work and want ready-to-use code solutions to cover key tasks in machine learning and the deep learning domain, then this book is what you need.

Familiarity with Python programming and machine learning concepts will be useful. Leverage deep learning to create powerful image processing apps with TensorFlow 2.0 and Keras Deep Learning in Healthcare Learning TensorFlow

Blue Book for World Internet Conference, Translated by CCTB Translation Service Computational Intelligence and Image Processing in Medical Applications High-Performance Big Data Computing Mastering TensorFlow 1.x

Comprehensive recipes to give you valuable insights on Transformers, Reinforcement Learning, and more Key Features Deep Learning solutions from Kaggle Masters and Google Developer Experts Get to grips with the fundamentals including variables, matrices, and data sources Learn advanced techniques to make your algorithms faster and more accurate Book Description The independent recipes in Machine Learning Using TensorFlow Cookbook will teach you how to perform complex data computations and gain valuable insights into your data. Dive into recipes on training models, model evaluation, sentiment analysis, regression analysis, artificial neural networks, and deep learning — each using Google's machine learning library, TensorFlow. This cookbook covers the fundamentals of the TensorFlow library, including variables, matrices, and various data sources. You'll discover real-world implementations of Keras and TensorFlow and learn how to use estimators to train linear models and boosted trees, both for classification and regression. Explore the practical applications of a variety of deep learning architectures, such as recurrent neural networks and Transformers, and see how they can be used to solve computer vision and natural language processing (NLP) problems. With the help of this book, you will be proficient in using TensorFlow, understand deep learning from the basics, and be able to implement machine learning algorithms in real-world scenarios. What you will learn Take TensorFlow into production Implement and fine-tune Transformer models for various NLP tasks Apply reinforcement learning algorithms using the TP-Agent framework Understand linear regression techniques and use Estimators to train linear models Execute neural networks and improve predictions on neural network tasks Build a recommender system through practical recipes Who this book is for If you are a data scientist or a machine learning engineer, and you want to skip detailed theoretical explanations in favor of building production-ready machine learning models using TensorFlow, this book is for you. Basic familiarity with Python, linear algebra, statistics, and machine learning is necessary to make the most out of this book.

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, Yehzekel 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

This book provides a comprehensive overview of deep learning (DL) in medical and healthcare applications, including the fundamentals and current advances in medical image analysis, state-of-the-art DL methods for medical image analysis and real-world, deep learning-based clinical computer-aided diagnosis systems. Deep learning (DL) is one of the key techniques of artificial intelligence (AI) and today plays an important role in numerous academic and industrial areas. DL involves using a neural network with many layers (deep structure) between input and output, and its main advantage is that it can automatically learn data-driven, highly representative and hierarchical features and perform feature extraction and classification on one network. DL can be used to model or simulate an intelligent system or process using annotated training data. Recently, DL has become widely used in medical applications, such as anatomic modelling, tumour detection, disease classification, computer-aided diagnosis and surgical planning. This book is intended for computer science and engineering students and researchers, medical professionals and anyone interested in using DL techniques.

The fundamental mathematical tools needed to understand machine learning include linear algebra, analytic geometry, matrix decompositions, vector calculus, optimization, probability and statistics. These topics are traditionally taught in disparate courses, making it hard for data science or computer science students, or professionals, to efficiently learn the mathematics. This self-contained textbook bridges the gap between mathematical and machine learning texts, introducing the mathematical concepts with a minimum of prerequisites. It uses these concepts to derive four central machine learning methods: linear regression, principal component analysis, Gaussian mixture models and support vector machines. For students and others with a mathematical background, these derivations provide a starting point to machine learning texts. For those learning the mathematics for the first time, the methods help build intuition and practical experience with applying mathematical concepts. Every chapter includes worked examples and exercises to test understanding. Programming tutorials are offered on the book's web site.

Develop Deep Learning Models for your Natural Language Problems Build Intelligent Language Applications Using Deep Learning Introduction to Deep Learning Reinforcement Learning

Create powerful machine learning algorithms with TensorFlow Embedded Deep Learning High Performance Computing for Big Data

Dependency-based methods for syntactic parsing have become increasingly popular in natural language processing in recent years. This book gives a thorough introduction to the methods that are most widely used today. After an introduction to dependency grammar and dependency parsing, followed by a formal characterization of the dependency parsing problem, the book surveys the three major classes of parsing models that are in current use: transition-based, graph-based, and grammar-based models. It continues with a chapter on evaluation and one on the comparison of different methods, and it closes with a few words on current trends and future prospects of dependency parsing. The book presupposes a knowledge of basic concepts in linguistics and computer science, as well as some knowledge of parsing methods for constituency-based representations. Table of Contents: Introduction / Dependency Parsing / Transition-Based Parsing / Graph-Based Parsing / Grammar-Based Parsing / Evaluation / Comparison / Final Thoughts

This book systematically reviews world internet development over the past 50 years, and comprehensively discusses the great contributions it has made to economic and social advances. Further, it describes the development, status and trends related to the Internet in major countries around the globe in 2019, and provides an in-depth analysis of the latest conditions, dynamics and development trends in key areas, including information infrastructure, information technology, digital economy, digital media, cyberspace security, and international cyberspace governance. Moreover, the book further modifies and enhances the Global Internet Development Index System in order to better show the Internet development strengths and advantages in various countries, and to reflect the global development trends more comprehensively, accurately and objectively. This book reviews the significant developments and summarizes the lessons learned as well as the future challenges. From a global perspective, it offers a vision of building a community with a shared future in cyberspace based on the new concepts, new ideas and new achievements of various countries participating in cyberspace development and construction. As such it is a valuable reference resource for anyone working in Internet related fields, such as those in government departments, internet enterprises, scientific research institutions, colleges and universities wanting to fully understand world Internet development.

Are all film stars linked to Kevin Bacon? Why do the stock markets rise and fall sharply on the strength of a vague rumour? How does gossip spread so quickly? Are we all related through six degrees of separation? There is a growing awareness of the complex networks that pervade modern society. We see them in the rapid growth of the Internet, the ease of global communication, the swift spread of news and information, and in the way epidemics and financial crisis develop with startling speed and intensity. This introductory book on the new science of networks takes an interdisciplinary approach, using how economics, sociology, computing, information science and applied mathematics to address fundamental questions about the links that connect us, and the ways that our decisions can have consequences for others.

Deep learning methods for achieving state-of-the-art results on challenging machine learning problems such as describing photos and translating text from one language to another. In this new laser-focused Ebook, finally cut through the math, research papers and patchwork descriptions about natural language processing. Using clear explanations, standard Python libraries and step-by-step tutorial lessons you will discover what natural language processing is, the promise of deep learning in the field, how to clean and prepare text data for modeling, and how to develop deep learning models for your own natural language processing projects.

Paradigms and Applications Methodologies and Applications

Algorithms, Architectures and Circuits for Always-on Neural Network Processing Deep Learning with JavaScript

Deep Learning State-of-the-Art

A Mathematical Approach to Advanced Artificial Intelligence in Python

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

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

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

Summary Deep learning has transformed the fields of computer vision, image processing, and natural language applications. Thanks to TensorFlow.js, now JavaScript developers can build deep learning apps without relying on Python or R. Deep Learning with JavaScript shows developers how they can bring DL technology to the web. Written by the main authors of the TensorFlow library, this new book provides fascinating use cases and in-depth instruction for deep learning apps in JavaScript in your browser or on Node. Foreword by Nikhil Thorat and Daniel Smilkov.

About the technology Running deep learning applications in the browser or on Node-based backends opens up exciting possibilities for smart web applications. With the TensorFlow.js library, you build and train deep learning models with JavaScript. Offering uncompromising production-quality scalability, modularity, and responsiveness, TensorFlow.js really shines for its portability. Its models run anywhere JavaScript runs, pushing ML farther up the application stack. About the book In Deep Learning with JavaScript, you'll learn to use TensorFlow.js to build deep learning models that run directly in the browser. This fast-paced book, written by Google engineers, is practical, engaging, and easy to follow. Through diverse examples featuring text analysis, speech processing, image recognition, and self-learning game AI, you'll master all the basics of deep learning and explore advanced concepts, like retraining existing models for transfer learning and image generation. What's inside - Image and language processing in the browser - Tuning ML models with client-side data - Text and image creation with generative deep learning - Source code samples to test and modify About the reader For JavaScript programmers interested in deep learning. About the author Shangqing Cai, Stanley Bilechki and Eric D. Nielsen are software engineers with experience on the Google Brain team, and were crucial to the development of the high-level API in TensorFlow.js. This book is based in part on the classic, Deep Learning with Python by François Chollet. TOC: PART 1 - MOTIVATION AND BASIC CONCEPTS 1 • Deep learning and JavaScript PART 2 - A GENTLE INTRODUCTION TO TENSORFLOW.JS 2 • Getting started: Simple linear regression in TensorFlow.js 3 • Adding nonlinearity: Beyond weighted sums 4 • Recognizing images and sounds using convnets 5 • Transfer learning: Reusing pretrained neural networks PART 3 - ADVANCED DEEP LEARNING WITH TENSORFLOW.JS 6 • Working with data 7 • Visualizing data and models 8 • Underfitting, overfitting, and the universal workflow of machine learning 9 • Deep learning for sequences and text 10 • Generative deep learning 11 • Basics of deep reinforcement learning PART 4 - SUMMARY AND CLOSING WORDS 12 • Testing, optimizing, and deploying models 13 • Summary, conclusions, and beyond

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

Create learning experiences that transform not only learning, but life itself. Learn about, improve, and expand your world of learning. This hands-on companion to the runaway best-seller, Deep Learning: Engage the World Change the World, provides an essential roadmap for building capacity in teachers, schools, districts, and systems to design deep learning, measure progress, and assess conditions needed to activate and sustain innovation. Loaded with tips, tools, protocols, and real-world examples, the easy-to-use guide has everything educators need to construct and drive meaningful deep learning experiences that give purpose, unleash student potential, and prepare students to become problem-solving change agents in a global society.

Recurrent Neural Networks with Python Quick Start Guide Machine Learning with TensorFlow, Second Edition

Pro Deep Learning with TensorFlow Tools for Engagement

Mining of Massive Datasets Graphical Models, Exponential Families, and Variational Inference

Dependency Parsing Deep Learning has already achieved remarkable results in many fields. Now it's making waves throughout the sciences broadly and the life sciences in particular. This practical book teaches developers and scientists how to use deep learning for genomics, chemistry, biophysics, microscopy, medical analysis, and other fields. Ideal for practicing developers and scientists ready to apply their skills to scientific applications such as biology, genetics, and drug discovery, this book introduces several deep network primitives. You'll follow a case study on the problem of designing new therapeutics that ties together physics, chemistry, biology, and medicine—an example that represents one of science's greatest challenges.

Learn the basics of performing machine learning on molecular data Understand why deep learning is a powerful tool for genetics and genomics Apply deep learning to understand biophysical systems Get a brief introduction to machine learning with DeepChem Use deep learning to analyze microscopic images Analyze medical scans using deep learning techniques About various neural network architectures and generative adversarial networks Interpret what your models are doing and how it's working Learn how to solve challenging machine learning problems with TensorFlow. Google's revolutionary new software library for deep learning. If you have some background in basic linear algebra and calculus, this practical book introduces machine-learning fundamentals by showing you how to design systems capable of detecting objects in images, understanding text, analyzing video, and predicting the properties of potential medicines. TensorFlow for Deep Learning teaches concepts through practical examples and helps you build knowledge of deep learning foundations from the ground up. It's ideal for practicing developers with experience designing software systems, and useful for scientists and other professionals familiar with scripting but not necessarily with designing learning algorithms. Learn TensorFlow fundamentals, including how to perform basic computation Build simple learning systems to understand their mathematical foundations Dive into fully connected deep networks used in thousands of applications Turn prototypes into high-quality models with hyperparameter optimization Process images with convolutional neural networks Handle natural language datasets with recurrent neural networks Use reinforcement learning to solve games such as tic-tac-toe Train deep networks with hardware including GPUs and tensor processing units

Now in its second edition, this book focuses on practical algorithms for mining data from even the largest datasets. Learn how to develop intelligent applications with sequential learning and apply modern methods for language modeling with neural network architectures for deep learning with Python's most popular TensorFlow framework. Key Features Train and deploy Recurrent Neural Networks using the popular TensorFlow library Apply long short-term memory units Expand your skills in complex neural network and deep learning topics Book Description Developers struggle to find an easy-to-follow learning resource for implementing Recurrent Neural Network (RNN) models. RNNs are the state-of-the-art model in deep learning for dealing with sequential data. From language translation to generating captions for an image, RNNs are used to continuously improve results. This book will teach you the fundamentals of RNNs, with example applications in Python and the TensorFlow library. The examples are accompanied by the right combination of theoretical knowledge and real-world implementations of concepts to build a solid foundation of neural network modeling. Your journey starts with the simplest RNN model, where you can grasp the fundamentals. The book then builds on this by proposing more advanced and complex algorithms. We use them to explain how a typical state-of-the-art RNN model works. From generating text to building a language translator, we show how some of today's most powerful AI applications work under the hood. After that, we turn to more advanced topics, along with developing skills in this exciting field. What you will learn Use TensorFlow to build RNN models Use the correct RNN architecture for a particular machine learning task Collect and clean the training data for your models Use the correct Python libraries for any task during the building phase of your model Optimize your model for higher accuracy Identify the differences between multiple models and how you can substitute them Learn the core deep learning fundamentals applicable to any machine learning model Who this book is for This book is for Machine Learning engineers and data scientists who want to learn about Recurrent Neural Network models with practical use-cases. Exposure to Python programming is required. Previous experience with TensorFlow will be helpful, but not mandatory.

Sequential learning and language modeling with TensorFlow Mathematics for Machine Learning

Machine Learning Using TensorFlow Cookbook Biomedical Signal Processing for Healthcare Applications

From Linear Regression to Reinforcement Learning Deep Learning for Natural Language Processing

Advanced machine learning and deep learning concepts using TensorFlow 1.x and Keras Updated with new code, new projects, and new chapters. Machine Learning with TensorFlow, Second Edition gives readers a solid foundation in machine-learning concepts and the TensorFlow library. Summary Updated with new code, new projects, and new chapters. Machine Learning with TensorFlow, Second Edition gives readers a solid foundation in machine-learning concepts and the TensorFlow library. Written by KAS O. OPL Demirel and Principal Engineer at Microsoft, this book is available for downloadable. Jupiter notebooks for a hands-on experience coding TensorFlow with Python. New and revised content expands coverage of core machine learning algorithms, and advancements in neural networks such as VGG-Face face identification classifiers and deep speech classifiers. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the technology Supercharge your data analysis with machine learning! ML algorithms automatically improve as they process data, so results get better over time. You don't have to be a mathematician to use ML. Tools like Google's TensorFlow library help with complex calculations so you can focus on getting the answers you need. About the book Machine Learning with TensorFlow, Second Edition is a fully revised guide to building machine learning models using Python and TensorFlow. You'll apply core ML concepts to real-world challenges, such as sentiment analysis, text classification, and image recognition. Hands-on examples illustrate neural network techniques for deep speech processing, facial identification, and auto-encoding with CIFAR-10. What's inside Machine Learning with TensorFlow Choosing the best ML approaches Visualizing algorithms with TensorBoard Sharing results with collaborators Running models in Docker. About the reader Requires intermediate Python skills and knowledge of general algebraic concepts like vectors and matrices. Examples use the super-stable 1.15.x branch of TensorFlow and TensorFlow 2.x. About the author Chris Mattmann is the Division Manager of the Artificial Intelligence, Analytics, and Innovation Organization at NASA Jet Propulsion Lab. The first edition of this book was written by Nishant Shukla with Kenneth Frickles. Table of Contents PART 1 - YOUR MACHINE LEARNING RIG 1 A machine-learning odyssey 2 TensorFlow essentials PART 2 - CORE LEARNING ALGORITHMS 3 Linear regression and beyond 4 Using regression for call-center volume prediction 5 A gentle introduction to classification 6 Sentiment classification: Large movie-review dataset 7 Automatically clustering data 8 Inferring user activity from Android accelerometer data 9 Hidden Markov models 10 Part-of-speech tagging and word-sense disambiguation PART 3 - THE NEURAL NETWORK PARADIGM 11 A peek into autoencoders 12 Applying autoencoders: The CIFAR-10 image dataset 13 Reinforcement learning 14 Convolutional neural networks 15 Building a real-world CNN: VGG-Face and VGG-Face Lite 16 Recurrent neural networks 17 LSTMs and automatic speech recognition 18 Sequence-to-sequence models for chatbots 19 Utility landscapes

Graph-structured data is ubiquitous throughout the natural and social sciences, from telecommunication networks to quantum chemistry. Building relational inductive biases into deep learning architectures is crucial for creating systems that can learn, reason, and generalize from this kind of data. Recent years have seen a surge in research on graph representation learning, including techniques for deep graph embeddings, generalizations of convolutional neural networks to graph-structured data, and neural message-passing approaches inspired by belief propagation. These advances in graph representation learning have led to new state-of-the-art results in numerous domains, including chemical synthesis, 3D vision, recommender systems, question answering, and social network analysis. This book provides a synthesis and overview of graph representation learning. It begins with a discussion of the goals of graph representation learning as well as key methodological foundations in graph theory and network analysis. Following this, the book introduces and reviews methods for learning node embeddings, including random-walk-based methods and applications to knowledge graphs. It then provides a technical synthesis and introduction to the highly successful graph neural network (GNN) formalism, which has become a dominant and fast-growing paradigm for deep learning with graph data. The book concludes with a synthesis of recent advancements in deep generative models for graphs—a nascent but quickly growing subset of graph representation learning. This book covers algorithmic and hardware implementation techniques to enable embedded deep learning. The authors describe synergistic design approaches on the application-, algorithm-, computer architecture-, and circuit-level that will help in achieving the goal of reducing the computational cost of deep learning algorithms. The impact of these techniques is displayed in four silicon prototypes for embedded deep learning. Gives a wide overview of a series of effective solutions for energy-efficient neural networks on battery constrained wearable devices; Discusses the optimization of neural networks for embedded deployment on all levels of the design hierarchy—applications, algorithms, hardware architectures, and circuits—supported by real silicon prototypes; Elaborates on how to design efficient Convolutional Neural Network processors, exploiting parallelism and data-rate, and low-precision computations; Supports the introduced theory and design concepts by four real silicon prototypes. The physical realization's implementation and achieved performances are discussed elaborately to illustrate and highlight the introduced cross-layer design concepts.

Summary Deep Learning with Python introduces the field of deep learning using the Python language and the powerful Keras library. Written by Keras creator and Google AI researcher François Chollet, this book builds your understanding through intuitive explanations and practical examples. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology Machine learning has made remarkable progress in recent years. We went from near-unspeakable speech and image recognition, to near-human accuracy. We went from machines that couldn't beat a serious Go player, to defeating a world champion. Behind this progress is deep learning—a combination of engineering advances, best practices, and theory that enables a wealth of previously impossible smart applications. About the Book Deep Learning with Python introduces the field of deep learning using the Python language and the powerful Keras library. Written by Keras creator and Google AI researcher François Chollet, this book builds your understanding through intuitive explanations and practical examples. You'll explore challenging concepts and practice with applications in computer vision, natural-language processing, and generative models. By the time you finish, you'll have the knowledge and hands-on skills to apply deep learning in your own projects. What's Inside Deep learning from First principles Setting up your own deep-learning environment Image-classification models Deep learning for text and sequences Neural style transfer, text generation, and image generation About the Reader Readers need intermediate Python skills. No previous experience with Keras, TensorFlow, or machine learning is required. About the Author François Chollet works on deep learning at Google in Mountain View, CA. He is the creator of the Keras deep-learning library, as well as a contributor to the TensorFlow machine-learning framework. He also does deep-learning research, with a focus on computer vision and the application of machine learning to formal reasoning. His papers have been published at major conferences in the field, including the Conference on Computer Vision and Pattern Recognition (CVPR), the Conference and Workshop on Neural Information Processing Systems (NIPS), the International Conference on Learning Representations (ICLR), and others. Table of Contents PART 1 - FUNDAMENTALS OF DEEP LEARNING What is deep learning? Before we begin: the mathematical building blocks of neural networks Fundamentals of machine learning PART 2 - DEEP LEARNING IN PRACTICE: Deep learning for computer vision Deep learning for text and sequences

Advanced deep-learning best practices Generative deep learning Conclusions appendix A - Installing Keras and its dependencies on Ubuntu appendix B - Running Jupiter notebooks on an EC2 GPU instance

A Guide to Building Deep Learning Systems Deep Learning for Coders with Fastai and PyTorch

Deep Learning for NLP and Speech Recognition Python Machine Learning Cookbook

Speech & Language Processing The A.I. Marketer

Over 90 unique recipes to solve artificial-intelligence driven problems with Python

A practical guide to building high performance systems for object detection, segmentation, video processing, smartphone applications, and more Key FeaturesDiscover how to build, train, and serve your own deep neural networks with TensorFlow 2 and KerasApply modern solutions to a wide range of applications such as object detection and video analysisLearn how to run your models on mobile devices and web pages and improve their performanceBook Description Computer vision solutions are becoming increasingly common, making their way into fields such as health, automobile, social media, and robotics. This book will help you explore TensorFlow 2, the brand new version of Google's open source framework for machine learning. You will understand how to benefit from using convolutional neural networks (CNNs) for visual tasks. Hands-On Computer Vision with TensorFlow 2 starts with the fundamentals of computer vision and deep learning, teaching you how to build a neural network from scratch. You will discover the features that have made TensorFlow the most widely used AI library, along with its intuitive Keras interface. You'll then move on to building, training, and deploying CNNs efficiently. Complete with concrete code examples, the book demonstrates how to classify images with modern solutions, such as Inception and ResNet, and extract specific content using You Only Look Once (YOLO), Mask R-CNN, and U-Net. You will also build generative adversarial networks (GANs) and variational autoencoders (VAEs) to create and edit images, and long short-term memory networks (LSTMs) to analyze videos. In the process, you will acquire advanced insights into transfer learning, data augmentation, domain adaptation, and mobile and web deployment, among other key concepts. By the end of the book, you will have both the theoretical understanding and practical skills to solve advanced computer vision problems with TensorFlow 2.0. What you will learnCreate your own neural networks from scratchClassify images with modern architectures including Inception and ResNetDetect and segment objects in images with YOLO, Mask R-CNN, and U-NetTackle problems faced when developing self-driving cars and facial emotion recognition systemsBoost your application's performance with transfer learning, GANs, and domain adaptationUse recurrent neural networks (RNNs) for video analysisOptimize and deploy your networks on mobile devices and in the browserWho this book is for If you're new to deep learning and have some background in Python programming and image processing, like reading/writing image files and editing pixels, this book is for you. Even if you're an expert curious about the new TensorFlow 2 features, you'll find this book useful. While some theoretical concepts require knowledge of algebra and calculus, the book covers concrete examples focused on practical applications such as visual recognition for self-driving cars and smartphone apps.

An in-depth overview of an emerging field that brings together high-performance computing, big data processing, and deep lEarning. Over the last decade, the exponential explosion of data known as big data has changed the way we understand and harness the power of data. The emerging field of high-performance big data computing, which brings together high-performance computing (HPC), big data processing, and deep learning, aims to meet the challenges posed by large-scale data processing. This book offers an in-depth overview of high-performance big data computing and the associated technical issues, approaches, and solutions. The book covers basic concepts and necessary background knowledge, including data processing frameworks, storage systems, and hardware capabilities; offers a detailed discussion of technical issues in accelerating big data computing in terms of computation, communication, memory and storage, codesign, workload characterization and benchmarking, and system deployment and management; and surveys benchmarks and workloads for evaluating big data middleware systems. It presents a detailed discussion of big data computing systems and applications with high-performance networking, computing, and storage technologies, including state-of-the-art designs for data processing and storage systems. Finally, the book considers some advanced research topics in high-performance big data computing, including designing high-performance deep learning over big data (DLoBD) stacks and HPC cloud technologies.

The core of this paper is a general set of variational principles for the problems of computing marginal probabilities and modes, applicable to multivariate statistical models in the exponential family.

Natural Language Processing (NLP) provides boundless opportunities for solving problems in artificial intelligence, making products such as Amazon Alexa and Google Translate possible. If you're a developer or data scientist new to NLP and deep learning, this practical guide shows you how to apply these methods using PyTorch, a Python-based deep learning library. Authors Delp Rao and Brian McMahon provide you with a solid grounding in NLP and deep learning algorithms and demonstrate how to use PyTorch to build applications involving rich representations of text specific to the problems you face. Each chapter includes several code examples and illustrations. Explore computational graphs and the supervised learning paradigm Master the basics of the PyTorch optimized tensor manipulation library Get an overview of traditional NLP concepts and methods Learn the basic ideas involved in building neural networks Use embeddings to represent words, sentences, documents, and other features Explore sequence prediction and generate sequence-to-sequence models Learn design patterns for building production NLP systems

Over 100 recipes to progress from smart data analytics to deep learning using real-world datasets, 2nd Edition

Graph Representation Learning

10 real-world projects on computer vision, machine translation, chatbots, and reinforcement learning

Deep Learning Systems

Neural networks in TensorFlow.js

Dive Into Deep Learning

TensorFlow for Deep Learning

Take the next step in implementing various common and not-so-common neural networks with TensorFlow 1.x About This Book Skill up and implement tricky neural networks using Google's TensorFlow 1.x An easy-to-follow guide that lets you explore reinforcement learning, GANs, autoencoders, multilayer perceptrons and more. Hands-on recipes to work with Tensorflow on desktop, mobile, and cloud environment Who This Book Is For This book is intended for data analysts, data scientists, machine learning practitioners and deep learning enthusiasts who want to perform deep learning tasks on a regular basis and are looking for a handy guide they can refer to. People who are slightly familiar with neural networks, and now want to gain expertise in working with different types of neur

networks and datasets, will find this book quite useful. What You Will Learn Install TensorFlow and use it for CPU and GPU operations Implement DNNs and apply them to solve different AI-driven problems. Leverage different data sets such as MNIST, CIFAR-10, and YouTube8m with TensorFlow and learn how to access and use them in your code. Use TensorBoard to understand neural network architectures, optimize the learning process, and peek inside the neural network black box. Use different regression techniques for prediction and classification problems Build single and multilayer perceptrons in TensorFlow Implement CNN and RNN in TensorFlow, and use it to solve real-world use cases. Learn how restricted Boltzmann Machines can be used to recommend movies. Understand the implementation of Autoencoders and deep belief networks, and use them for emotion detection. Master the different reinforcement learning methods to implement game playing agents, GANs and their implementation using TensorFlow. In Detail Deep neural networks (DNNs) have achieved a lot of success in the field of computer vision, speech recognition, and natural language processing. Th

entire world is filled with excitement about how deep networks are revolutionizing artificial intelligence. This exciting recipe-based guide will take you from the realm of DNN theory to implementing them practically to solve the real-life problems in artificial intelligence domain. In this book, you will learn how to efficiently use TensorFlow, Google's open source framework for deep learning. You will implement different deep learning networks such as Convolutional Neural Networks (CNNs), Recurrent Neural Networks (RNNs), Deep Q-learning Networks (DQNs), and Generative Adversarial Networks (GANs) with easy to follow independent recipes. You will learn how to make Keras as backend with TensorFlow. With a problem-solution approach, you will understand how to implement different deep neural architectures to carry out complex tasks at work. You will learn the performance of different DNNs on some popularly used data sets such as MNIST, CIFAR-10, YouTube8m, and more. You will not only learn about the different mobile and embedded platforms supported by TensorFlow but also how to set up cloud platforms for deep learning applications. Get a sneak peek of TPU architecture and how they will affect DNN future. By using crisp, no-nonsense recipes, you will become an expert in implementing deep learning techniques in growing real-world applications and research areas such as reinforcement learning, GANs, autoencoders and more. Style and approach This book consists of hands-on recipes where you'll deal with real-world problems. You'll execute a serie

of tasks as you walk through data mining challenges using TensorFlow 1.x. Your one-stop solution for common and not-so-common pain points, this is a book that you must have on the shelf. This textbook explains Deep Learning Architecture, with applications to various NLP Tasks, including Document Classification, Machine Translation, Language Modeling, and Speech Recognition. With the widespread adoption of deep learning, natural language processing (NLP) and speech applications in many areas (including Finance, Healthcare, and Government) there is a growing need for one comprehensive resource that maps deep learning techniques to NLP and speech and provides insights into using the tools and libraries for real-world applications. Deep Learning for NLP and Speech Recognition explains recent deep learning methods applicable to NLP and speech, provides state-of-the-art approaches, and offers real-world case studies with code to provide hands-on experience.

Many books focus on deep learning theory or deep learning for NLP-specific tasks while others are cookbooks for tools and libraries, but the constant flux of new algorithms, tools, frameworks, and libraries in a rapidly evolving landscape means that there are few available texts that offer the material in this book. The book is organized into three parts, aligning to different groups of readers and expertise. The three parts are: Machine Learning, NLP, and Speech Introduction The first part has three chapters that introduce readers to the fields of NLP, speech recognition, deep learning and machine learning with basic theory and hands-on case studies using Python-based tools and libraries. Deep Learning Basics The five chapters in the second part introduce deep learning and various topics that are crucial for speech and text processing, including word embeddings, convolutional neural networks, recurrent neural networks and speech recognition basics. Theory, practical tips, state-of-the-art methods, experimentations and analysis in using the methods discussed in theory on real-world tasks. Advanced Deep Learning Techniques for Text and Speech The third part has five chapters that discuss the latest and cutting-edge research in the areas of deep learning that intersect with NLP and speech. Topics including attention mechanisms, memory augmented networks, transfer learning, multi-task learning, domain adaptation, reinforcement learning, and end-to-end deep learning for speech recognition are covered using case studies.

Build, scale, and deploy deep neural network models using the star libraries in Python Key Features Delve into advanced machine learning and deep learning use cases using TensorFlow and Keras Build, deploy, and scale end-to-end deep neural network models in a production environment Learn to deploy TensorFlow on mobile, and distributed TensorFlow on GPU, Clusters, and Kubernetes Book Description TensorFlow is the most popular numerical computation library built from the ground up for distributed, cloud, and mobile environments. TensorFlow represents the data as tensors and the computation as graphs. This book is a comprehensive guide that lets you explore the advanced features of TensorFlow 1.x. Gain insight into TensorFlow Core, Keras, TF Estimators, TFLearn, TFSlim, Pretty Tensor, and Sonnet. Leverage the power of TensorFlow and Keras to build deep learning models, using concepts such as transfer learning, generative adversarial networks, and deep reinforcement learning. Throughout the book, you will obtain hands-on experience with varied datasets, such as MNIST, CIFAR-10, PTB, text8, and COCO-Images. You will learn the advanced features of TensorFlow1.x, such as distributed TensorFlow with TF Clusters, deploy production models with TensorFlow Serving, and build and deploy TensorFlow models for mobile and embedded devices on Android and iOS platforms. You will see how to call TensorFlow and Keras API within the R statistical software, and learn the required techniques for debugging when the TensorFlow API-based code does not work as expected. The book helps you obtain in-depth knowledge of TensorFlow, making you the go-to person for solving artificial intelligence problems. By the end of this guide, you will have mastered the offerings of TensorFlow and Keras, and gained the skills you need to build smarter, faster, and efficient machine learning and deep learning systems. What you will learn Master advanced concepts of deep learning such as transfer learning, reinforcement learning, generative models and more, using TensorFlow and Keras Perform supervised (classification and regression) and unsupervised (clustering) learning to solve machine learning tasks Build end-to-end deep learning (CNN, RNN, and Autoencoders) models with TensorFlow Scale and deploy production models with distributed and high-performance computing on GPU and clusters Build TensorFlow models to work with multilayer perceptrons using Keras, TFLearn, and R Learn the functionalities of smart apps by building and deploying TensorFlow models on iOS and Android devices Supercharge TensorFlow with distributed training and deployment on Kubernetes and TensorFlow Clusters Who this book is for

This book is for data scientists, machine learning engineers, artificial intelligence engineers, and for all TensorFlow users who wish to upgrade their TensorFlow knowledge and work on various machine learning and deep learning problems. If you are looking for an easy-to-follow guide that underlines the intricacies and complex use cases of machine learning, you will find this book extremely useful. Some basic understanding of TensorFlow is required to get the most out of the book.

This book examines the use of biomedical signal processing—EEG, EMG, and ECG—in analyzing and diagnosing various medical conditions, particularly diseases related to the heart and brain. In combination with machine learning tools and other optimization methods, the analysis of biomedical signals greatly benefits the healthcare sector by improving patient outcomes through early, reliable detection. The discussion of these modalities promotes better understanding, analysis, and application of biomedical signal processing for specific diseases. The major highlights of Biomedical Signal Processing for Healthcare Applications include biomedical signals, acquisition of signals, pre-processing and analysis, post-processing and classification of the signals, and application of analysis and classification for the diagnosis of brain- and heart-related diseases. Emphasis is given to brain and heart signals because incomplete interpretations are made by physicians of these aspects in several situations, and these partial interpretations lead to major complications. FEATURES Examines modeling and acquisition of biomedical signals of different disorders Discusses CAD-based analysis of diagnosis useful for healthcare Includes all important modalities of biomedical signals, such as EEG, EMG, MEG, ECG, and PCG Includes case studies and research directions, including novel approaches used in advanced healthcare systems This book can be used by a wide range of users, including students, research scholars, faculty, and practitioners in the field of biomedical engineering and medical image analysis and diagnosis.

Networks, Crowds, and Markets

TensorFlow 1.x Deep Learning Cookbook

TensorFlow Machine Learning Cookbook

A Visual Approach

TensorFlow Deep Learning Projects

World Internet Development Report 2019

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.