

Tflearn Tensorflow Deep Learning Library

As recognized, adventure as capably as experience just about lesson, amusement, as with ease as accord can be gotten by just checking out a book **Tflearn Tensorflow Deep Learning Library** as well as it is not directly done, you could admit even more in relation to this life, in this area the world.

We give you this proper as well as easy mannerism to acquire those all. We offer Tflearn Tensorflow Deep Learning Library and numerous book collections from fictions to scientific research in any way. among them is this Tflearn Tensorflow Deep Learning Library that can be your partner.

Deep Learning With Python - Jason Brownlee 2016-05-13

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

Learning TensorFlow - Tom Hope 2017-08-09

Roughly inspired by the human brain, deep neural networks trained with large amounts of data can solve complex tasks with unprecedented accuracy. This practical book provides an end-to-end guide to TensorFlow, the leading open source software library that helps you build and train neural networks for computer vision, natural language processing (NLP), speech recognition, and general predictive analytics. Authors Tom Hope, Yehezkel Resheff, and Itay Lieder provide a hands-on approach to TensorFlow fundamentals for a broad technical audience—from data scientists and engineers to students and researchers. You'll begin by working through some basic examples in TensorFlow before diving deeper into topics such as neural network architectures, TensorBoard visualization, TensorFlow abstraction libraries, and multithreaded input pipelines. Once you finish this book, you'll know how to build and deploy production-ready deep learning systems in TensorFlow. Get up and running with TensorFlow, rapidly and painlessly Learn how to use TensorFlow to build deep learning models from the ground up Train popular deep learning models for computer vision and NLP Use extensive abstraction libraries to make development easier and faster Learn how to scale TensorFlow, and use clusters to distribute model training Deploy TensorFlow in a production setting

Applied Cloud Deep Semantic Recognition - Mehdi Roopaei 2018-04-09

This book provides a comprehensive overview of the research on anomaly detection with respect to context and situational awareness that aim to get a better understanding of how context information influences anomaly detection. In each chapter, it identifies advanced anomaly detection and key assumptions, which are used by the model to differentiate between normal and anomalous behavior. When applying a given model to a particular application, the assumptions can be used as guidelines to assess the effectiveness of the model in that domain. Each chapter provides an advanced deep content understanding and anomaly detection algorithm, and then shows how the proposed approach is deviating of the basic techniques. Further, for each chapter, it describes the advantages and disadvantages of the algorithm. The final chapters provide a discussion on the computational complexity of the models and graph computational frameworks such as Google Tensorflow and H2O because it is an important issue in real application domains. This book provides a better understanding of the different directions in which research has been done on deep semantic analysis and situational assessment using deep learning for anomalous detection, and how methods developed in one area can be applied in applications in other domains. This book seeks to provide both cyber analytics practitioners and researchers an up-to-date and advanced knowledge in cloud based frameworks for deep semantic analysis and advanced anomaly detection using cognitive and artificial intelligence (AI) models.

Data Science Programming All-in-One For Dummies - John Paul Mueller 2020-01-09

Your logical, linear guide to the fundamentals of data science programming Data science is exploding—in a good way—with a forecast of 1.7 megabytes of new information created every second for each human being on the planet by 2020 and 11.5 million job openings by 2026. It clearly pays dividends to be in the know. This friendly guide charts a path through the fundamentals of data science and then delves into the actual work: linear regression, logical regression, machine learning, neural networks, recommender engines, and cross-validation of

models. Data Science Programming All-In-One For Dummies is a compilation of the key data science, machine learning, and deep learning programming languages: Python and R. It helps you decide which programming languages are best for specific data science needs. It also gives you the guidelines to build your own projects to solve problems in real time. Get grounded: the ideal start for new data professionals What lies ahead: learn about specific areas that data is transforming Be meaningful: find out how to tell your data story See clearly: pick up the art of visualization Whether you're a beginning student or already mid-career, get your copy now and add even more meaning to your life—and everyone else's!

Innovative Computing - Chao-Tung Yang 2020-09-25

This book gathers peer-reviewed proceedings of the 3rd International Conference on Innovative Computing (IC 2020). This book aims to provide an open forum for discussing recent advances and emerging trends in information technology, science, and engineering. Themes within the scope of the conference include Communication Networks, Business Intelligence and Knowledge Management, Web Intelligence, and any related fields that depend on the development of information technology. The respective contributions presented here cover a wide range of topics, from databases and data mining, networking and communications, the web and Internet of Things, to embedded systems, soft computing, social network analysis, security and privacy, optical communication, and ubiquitous/pervasive computing. Readers such as students, researchers, and industry professionals in the fields of cloud computing, Internet of Things, machine learning, information security, multimedia systems, and information technology benefit from this comprehensive overview of the latest advances in information technology. The book can also benefit young investigators looking to start a new research program.

TensorFlow Machine Learning Projects - Ankit Jain 2018-11-30

Implement TensorFlow's offerings such as TensorBoard, TensorFlow.js, TensorFlow Probability, and TensorFlow Lite to build smart automation projects Key FeaturesUse machine learning and deep learning principles to build real-world projectsGet to grips with TensorFlow's impressive range of module offeringsImplement projects on GANs, reinforcement learning, and capsule networkBook Description TensorFlow has transformed the way machine learning is perceived. TensorFlow Machine Learning Projects teaches you how to exploit the benefits—simplicity, efficiency, and flexibility—of using TensorFlow in various real-world projects. With the help of this book, you'll not only learn how to build advanced projects using different datasets but also be able to tackle common challenges using a range of libraries from the TensorFlow ecosystem. To start with, you'll get to grips with using TensorFlow for machine learning projects; you'll explore a wide range of projects using TensorForest and TensorBoard for detecting exoplanets, TensorFlow.js for sentiment analysis, and TensorFlow Lite for digit classification. As you make your way through the book, you'll build projects in various real-world domains, incorporating natural language processing (NLP), the Gaussian process, autoencoders, recommender systems, and Bayesian neural networks, along with trending areas such as Generative Adversarial Networks (GANs), capsule networks, and reinforcement learning. You'll learn how to use the TensorFlow on Spark API and GPU-accelerated computing with TensorFlow to detect objects, followed by how to train and develop a recurrent neural network (RNN) model to generate book scripts. By the end of this book, you'll have gained the required expertise to build full-fledged machine learning projects at work. What you will learnUnderstand the TensorFlow ecosystem using various datasets and techniquesCreate recommendation systems for quality product recommendationsBuild projects using CNNs, NLP, and Bayesian neural networksPlay Pac-Man using deep reinforcement learningDeploy scalable TensorFlow-based machine learning systemsGenerate your own book script using RNNsWho this book is for

TensorFlow Machine Learning Projects is for you if you are a data analyst, data scientist, machine learning professional, or deep learning enthusiast with basic knowledge of TensorFlow. This book is also for you if you want to build end-to-end projects in the machine learning domain using supervised, unsupervised, and reinforcement learning techniques

[Deep Learning Systems](#) - Andres Rodriguez 2022-05-31

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 codesign 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 that 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 platform constraints when designing models. Performance engineers should support optimizations across diverse models, 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 stack 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.

[Cognitive Computing Recipes](#) - Adnan Masood 2019-03-27

Solve your AI and machine learning problems using complete and real-world code examples. Using a problem-solution approach, this book makes deep learning and machine learning accessible to everyday developers, by providing a combination of tools such as cognitive services APIs, machine learning platforms, and libraries. Along with an overview of the contemporary technology landscape, Machine Learning and Deep Learning with Cognitive Computing Recipes covers the business case for machine learning and deep learning. Covering topics such as digital assistants, computer vision, text analytics, speech, and robotics process automation this book offers a comprehensive toolkit that you can apply quickly and easily in your own projects. With its focus on Microsoft Cognitive Services offerings, you'll see recipes using multiple different environments including TensorFlow and CNTK to give you a broader perspective of the deep learning ecosystem. What You Will Learn Build production-ready solutions using Microsoft Cognitive Services APIs Apply deep learning using TensorFlow and Microsoft Cognitive Toolkit (CNTK) Solve enterprise problems in natural language processing and computer vision Discover the machine learning development life cycle - from formal problem definition to deployment at scale Who This Book Is For Software engineers and enterprise architects who wish to understand machine learning and deep learning by building applications and solving real-world business problems.

[Python: Advanced Guide to Artificial Intelligence](#) - Giuseppe Bonaccorso 2018-12-21

Demystify the complexity of machine learning techniques and create evolving, clever solutions to solve your problems Key Features Master supervised, unsupervised, and semi-supervised ML algorithms and their

implementation Build deep learning models for object detection, image classification, similarity learning, and more Build, deploy, and scale end-to-end deep neural network models in a production environment Book Description This Learning Path is your complete guide to quickly getting to grips with popular machine learning algorithms. You'll be introduced to the most widely used algorithms in supervised, unsupervised, and semi-supervised machine learning, and learn how to use them in the best possible manner. Ranging from Bayesian models to the MCMC algorithm to Hidden Markov models, this Learning Path will teach you how to extract features from your dataset and perform dimensionality reduction by making use of Python-based libraries. You'll bring the use of TensorFlow and Keras to build deep learning models, using concepts such as transfer learning, generative adversarial networks, and deep reinforcement learning. Next, you'll learn the advanced features of TensorFlow 1.x, such as distributed TensorFlow with TF clusters, deploy production models with TensorFlow Serving. You'll implement different techniques related to object classification, object detection, image segmentation, and more. By the end of this Learning Path, you'll have obtained in-depth knowledge of TensorFlow, making you the go-to person for solving artificial intelligence problems This Learning Path includes content from the following Packt products: Mastering Machine Learning Algorithms by Giuseppe Bonaccorso Mastering TensorFlow 1.x by Armando Fandango Deep Learning for Computer Vision by Rajalingappaa Shanmugamani What you will learn Explore how an ML model can be trained, optimized, and evaluated Work with Autoencoders and Generative Adversarial Networks Explore the most important Reinforcement Learning techniques Build end-to-end deep learning (CNN, RNN, and Autoencoders) models Who this book is for This Learning Path is for data scientists, machine learning engineers, artificial intelligence engineers who want to delve into complex machine learning algorithms, calibrate models, and improve the predictions of the trained model. You will encounter the advanced intricacies and complex use cases of deep learning and AI. A basic knowledge of programming in Python and some understanding of machine learning concepts are required to get the best out of this Learning Path.

[Deep Learning with Applications Using Python](#) - Navin Kumar Manaswi 2018-04-04

Explore deep learning applications, such as computer vision, speech recognition, and chatbots, using frameworks such as TensorFlow and Keras. This book helps you to ramp up your practical know-how in a short period of time and focuses you on the domain, models, and algorithms required for deep learning applications. Deep Learning with Applications Using Python covers topics such as chatbots, natural language processing, and face and object recognition. The goal is to equip you with the concepts, techniques, and algorithm implementations needed to create programs capable of performing deep learning. This book covers convolutional neural networks, recurrent neural networks, and multilayer perceptrons. It also discusses popular APIs such as IBM Watson, Microsoft Azure, and scikit-learn. What You Will Learn Work with various deep learning frameworks such as TensorFlow, Keras, and scikit-learn. Use face recognition and face detection capabilities Create speech-to-text and text-to-speech functionality Engage with chatbots using deep learning Who This Book Is For Data scientists and developers who want to adapt and build deep learning applications.

[Deep Learning with TensorFlow](#) - Giancarlo Zaccone 2017-04-24

Delve into neural networks, implement deep learning algorithms, and explore layers of data abstraction with the help of this comprehensive TensorFlow guide About This Book Learn how to implement advanced techniques in deep learning with Google's brainchild, TensorFlow Explore deep neural networks and layers of data abstraction with the help of this comprehensive guide Real-world contextualization through some deep learning problems concerning research and application Who This Book Is For The book is intended for a general audience of people interested in machine learning and machine intelligence. A rudimentary level of programming in one language is assumed, as is a basic familiarity with computer science techniques and technologies, including a basic awareness of computer hardware and algorithms. Some competence in mathematics is needed to the level of elementary linear algebra and calculus. What You Will Learn Learn about machine learning landscapes along with the historical development and progress of deep learning Learn about deep machine intelligence and GPU computing with the latest TensorFlow 1.x Access public datasets and utilize them using TensorFlow to load, process, and transform data Use TensorFlow on real-world datasets, including images, text, and more Learn how to evaluate the performance of your deep learning models Using deep learning for

scalable object detection and mobile computing Train machines quickly to learn from data by exploring reinforcement learning techniques Explore active areas of deep learning research and applications In Detail Deep learning is the step that comes after machine learning, and has more advanced implementations. Machine learning is not just for academics anymore, but is becoming a mainstream practice through wide adoption, and deep learning has taken the front seat. As a data scientist, if you want to explore data abstraction layers, this book will be your guide. This book shows how this can be exploited in the real world with complex raw data using TensorFlow 1.x. Throughout the book, you'll learn how to implement deep learning algorithms for machine learning systems and integrate them into your product offerings, including search, image recognition, and language processing. Additionally, you'll learn how to analyze and improve the performance of deep learning models. This can be done by comparing algorithms against benchmarks, along with machine intelligence, to learn from the information and determine ideal behaviors within a specific context. After finishing the book, you will be familiar with machine learning techniques, in particular the use of TensorFlow for deep learning, and will be ready to apply your knowledge to research or commercial projects. Style and approach This step-by-step guide will explore common, and not so common, deep neural networks and show how these can be exploited in the real world with complex raw data. With the help of practical examples, you will learn how to implement different types of neural nets to build smart applications related to text, speech, and image data processing.

[Proceedings of Second International Conference on Computing,](#)

[Communications, and Cyber-Security](#) - Pradeep Kumar Singh 2021-05-24

This book features selected research papers presented at the Second International Conference on Computing, Communications, and Cyber-Security (IC4S 2020), organized in Krishna Engineering College (KEC), Ghaziabad, India, along with Academic Associates; Southern Federal University, Russia; IAC Educational, India; and ITS Mohan Nagar, Ghaziabad, India during 3-4 October 2020. It includes innovative work from researchers, leading innovators, and professionals in the area of communication and network technologies, advanced computing technologies, data analytics and intelligent learning, the latest electrical and electronics trends, and security and privacy issues.

Fundamentals of Deep Learning - Nikhil Buduma 2017-05-25

With the reinvigoration of neural networks in the 2000s, deep learning has become an extremely active area of research, one that's paving the way for modern machine learning. In this practical book, author Nikhil Buduma provides examples and clear explanations to guide you through major concepts of this complicated field. Companies such as Google, Microsoft, and Facebook are actively growing in-house deep-learning teams. For the rest of us, however, deep learning is still a pretty complex and difficult subject to grasp. If you're familiar with Python, and have a background in calculus, along with a basic understanding of machine learning, this book will get you started. Examine the foundations of machine learning and neural networks Learn how to train feed-forward neural networks Use TensorFlow to implement your first neural network Manage problems that arise as you begin to make networks deeper Build neural networks that analyze complex images Perform effective dimensionality reduction using autoencoders Dive deep into sequence analysis to examine language Learn the fundamentals of reinforcement learning

[Mastering TensorFlow 1.x](#) - Armando Fandango 2018-01-22

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, TF Slim, 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.

Python Machine Learning - Ryan Turner 2020-04-18

Are you a novice programmer who wants to learn Python Machine Learning? Are you worried about how to translate what you already know into Python? This book will help you overcome those problems! As machines get ever more complex and perform more and more tasks to free up our time, so it is that new ideas are developed to help us continually improve their speed and abilities. One of these is Python and in Python Machine Learning: 3 books in 1 - The Ultimate Beginner's Guide to Learn Python Machine Learning Step by Step using Scikit-Learn and Tensorflow, you will discover information and advice on: Book 1 • What machine learning is • The history of machine learning • Approaches to machine learning • Support vector machines • Machine learning and neural networks • The Internet of Things (IoT) • The future of machine learning • And more... Book 2 • The principles surrounding Python • Different types of networks so you can choose what works best for you • Features of the system • Real world feature engineering • Understanding the techniques of semi-supervised learning • And more... Book 3 • How advanced tensorflow can be used • Neural network models and how to get the most from them • Machine learning with Generative Adversarial Networks • Translating images with cross domain GANs • TF clusters and how to use them • How to debug TF models • And more...

This book has been written specifically for beginners and the simple, step by step instructions and plain language make it an ideal place to start for anyone who has a passing interest in this fascinating subject. Python really is an amazing system and can provide you with endless possibilities when you start learning about it. Get a copy of Python Machine Learning today and see where the future lies.

[Internet of Things and Analytics for Agriculture, Volume 3](#) - Prasant Kumar Pattnaik 2021-11-10

The book discusses one of the major challenges in agriculture which is delivery of cultivate produce to the end consumers with best possible price and quality. Currently all over the world, it is found that around 50% of the farm produce never reaches the end consumer due to wastage and suboptimal prices. The authors present solutions to reduce the transport cost, predictability of prices on the past data analytics and the current market conditions, and number of middle hops and agents between the farmer and the end consumer using IoT-based solutions. Again, the demand by consumption of agricultural products could be predicted quantitatively; however, the variation of harvest and production by the change of farm's cultivated area, weather change, disease and insect damage, etc., could be difficult to be predicted, so that the supply and demand of agricultural products has not been controlled properly. To overcome, this edited book designed the IoT-based monitoring system to analyze crop environment and the method to improve the efficiency of decision making by analyzing harvest statistics. The book is also useful for academicians working in the areas of climate changes.

TensorFlow 1.x Deep Learning Cookbook - Antonio Gulli 2017-12-12
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 neural 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. The 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 series 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.

Computer Vision Projects with OpenCV and Python 3 - Matthew Rever 2018-12-28

Gain a working knowledge of advanced machine learning and explore Python's powerful tools for extracting data from images and videos Key Features Implement image classification and object detection using machine learning and deep learning Perform image classification, object detection, image segmentation, and other Computer Vision tasks Crisp content with a practical approach to solving real-world problems in Computer Vision Book Description Python is the ideal programming language for rapidly prototyping and developing production-grade codes for image processing and Computer Vision with its robust syntax and wealth of powerful libraries. This book will help you design and develop production-grade Computer Vision projects tackling real-world problems. With the help of this book, you will learn how to set up Anaconda and Python for the major OSes with cutting-edge third-party libraries for Computer Vision. You'll learn state-of-the-art techniques for classifying images, finding and identifying human postures, and detecting faces within videos. You will use powerful machine learning tools such as OpenCV, Dlib, and TensorFlow to build exciting projects such as classifying handwritten digits, detecting facial features, and much more.

The book also covers some advanced projects, such as reading text from license plates from real-world images using Google's Tesseract software, and tracking human body poses using DeeperCut within TensorFlow. By the end of this book, you will have the expertise required to build your own Computer Vision projects using Python and its associated libraries. What you will learn Install and run major Computer Vision packages within Python Apply powerful support vector machines for simple digit classification Understand deep learning with TensorFlow Build a deep learning classifier for general images Use LSTMs for automated image captioning Read text from real-world images Extract human pose data from images Who this book is for Python programmers and machine learning developers who wish to build exciting Computer Vision projects using the power of machine learning and OpenCV will find this book useful. The only prerequisite for this book is that you should have a sound knowledge of Python programming.

Hands-On Machine Learning with Scikit-Learn and TensorFlow - Aurélien Géron 2017-03-13

Through a series of recent breakthroughs, deep learning has boosted the entire field of machine learning. Now, even programmers who know close to nothing about this technology can use simple, efficient tools to implement programs capable of learning from data. This practical book shows you how. By using concrete examples, minimal theory, and two production-ready Python frameworks--scikit-learn and TensorFlow--author Aurélien Géron helps you gain an intuitive understanding of the concepts and tools for building intelligent systems. You'll learn a range of techniques, starting with simple linear regression and progressing to deep neural networks. With exercises in each chapter to help you apply what you've learned, all you need is programming experience to get started.

Sustainable Communication Networks and Application - P. Karuppusamy 2021-01-25

This book includes novel and state-of-the-art research discussions that articulate and report all research aspects, including theoretical and experimental prototypes and applications that incorporate sustainability into emerging applications. In recent years, sustainability and information and communication technologies (ICT) are highly intertwined, where sustainability resources and its management has attracted various researchers, stakeholders, and industrialists. The energy-efficient communication technologies have revolutionized the various smart applications like smart cities, healthcare, entertainment, and business. The book discusses and articulates emerging challenges in significantly reducing the energy consumption of communication systems and also explains development of a sustainable and energy-efficient mobile and wireless communication network. It includes best selected high-quality conference papers in different fields such as internet of things, cloud computing, data mining, artificial intelligence, machine learning, autonomous systems, deep learning, neural networks, renewable energy sources, sustainable wireless communication networks, QoS, network sustainability, and many other related areas.

Deep Learning, Machine Learning and IoT in Biomedical and Health Informatics - Sujata Dash 2022-02-11

Biomedical and Health Informatics is an important field that brings tremendous opportunities and helps address challenges due to an abundance of available biomedical data. This book examines and demonstrates state-of-the-art approaches for IoT and Machine Learning based biomedical and health related applications. This book aims to provide computational methods for accumulating, updating and changing knowledge in intelligent systems and particularly learning mechanisms that help us to induce knowledge from the data. It is helpful in cases where direct algorithmic solutions are unavailable, there is lack of formal models, or the knowledge about the application domain is inadequately defined. In the future IoT has the impending capability to change the way we work and live. These computing methods also play a significant role in design and optimization in diverse engineering disciplines. With the influence and the development of the IoT concept, the need for AI (artificial intelligence) techniques has become more significant than ever. The aim of these techniques is to accept imprecision, uncertainties and approximations to get a rapid solution. However, recent advancements in representation of intelligent IoT systems generate a more intelligent and robust system providing a human interpretable, low-cost, and approximate solution. Intelligent IoT systems have demonstrated great performance to a variety of areas including big data analytics, time series, biomedical and health informatics. This book will be very beneficial for the new researchers and practitioners working in the biomedical and healthcare fields to quickly know the best performing

methods. It will also be suitable for a wide range of readers who may not be scientists but who are also interested in the practice of such areas as medical image retrieval, brain image segmentation, among others. • Discusses deep learning, IoT, machine learning, and biomedical data analysis with broad coverage of basic scientific applications • Presents deep learning and the tremendous improvement in accuracy, robustness, and cross-language generalizability it has over conventional approaches • Discusses various techniques of IoT systems for healthcare data analytics • Provides state-of-the-art methods of deep learning, machine learning and IoT in biomedical and health informatics • Focuses more on the application of algorithms in various real life biomedical and engineering problems

Algorithms and Architectures for Parallel Processing - Jaideep Vaidya 2018-12-06

The four-volume set LNCS 11334-11337 constitutes the proceedings of the 18th International Conference on Algorithms and Architectures for Parallel Processing, ICA3PP 2018, held in Guangzhou, China, in November 2018. The 141 full and 50 short papers presented were carefully reviewed and selected from numerous submissions. The papers are organized in topical sections on Distributed and Parallel Computing; High Performance Computing; Big Data and Information Processing; Internet of Things and Cloud Computing; and Security and Privacy in Computing.

Smart Agricultural Services Using Deep Learning, Big Data, and IoT - Gupta, Amit Kumar 2020-10-30

The agricultural sector can benefit immensely from developments in the field of smart farming. However, this research area focuses on providing specific fixes to particular situations and falls short on implementing data-driven frameworks that provide large-scale benefits to the industry as a whole. Using deep learning can bring immense data and improve our understanding of various earth sciences and improve farm services to yield better crop production and profit. *Smart Agricultural Services Using Deep Learning, Big Data, and IoT* is an essential publication that focuses on the application of deep learning to agriculture. While highlighting a broad range of topics including crop models, cybersecurity, and sustainable agriculture, this book is ideally designed for engineers, programmers, software developers, agriculturalists, farmers, policymakers, researchers, academicians, and students.

Deep Learning with JavaScript - Stanley Bileschi 2020-01-24

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 Shangting Cai, Stanley Bileschi 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 of 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

Progress in Advanced Computing and Intelligent Engineering - Bibudhendu Pati 2018-12-17

The book gathers high-quality research papers presented at the International Conference on Advanced Computing and Intelligent Engineering (ICACIE 2017). It includes technical sections describing progress in the fields of advanced computing and intelligent engineering, and is primarily intended for postgraduate students and researchers working in Computer Science and Engineering. However, researchers working in Electronics will also find the book useful, as it addresses hardware technologies and next-gen communication technologies.

Deep Learning with Python - Francois Chollet 2017-11-30

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-unusable 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 Getting started with 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 Jupyter notebooks on an EC2 GPU instance

Applied Text Analysis with Python - Benjamin Bengfort 2018-06-11

From news and speeches to informal chatter on social media, natural language is one of the richest and most underutilized sources of data. Not only does it come in a constant stream, always changing and adapting in context; it also contains information that is not conveyed by traditional data sources. The key to unlocking natural language is through the creative application of text analytics. This practical book presents a data scientist's approach to building language-aware products with applied machine learning. You'll learn robust, repeatable, and scalable techniques for text analysis with Python, including contextual and linguistic feature engineering, vectorization, classification, topic modeling, entity resolution, graph analysis, and visual steering. By the end of the book, you'll be equipped with practical methods to solve any number of complex real-world problems. Preprocess and vectorize text

into high-dimensional feature representations Perform document classification and topic modeling Steer the model selection process with visual diagnostics Extract key phrases, named entities, and graph structures to reason about data in text Build a dialog framework to enable chatbots and language-driven interaction Use Spark to scale processing power and neural networks to scale model complexity *Machine Learning with TensorFlow, Second Edition* - Mattmann A. Chris 2021-02-02

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 NASA JPL Deputy CTO and Principal Data Scientist Chris Mattmann, all examples are accompanied by downloadable Jupyter 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 facial 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 Fricklas. 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 ad VGG-Face Lite 16 Recurrent neural networks 17 LSTMs and automatic speech recognition 18 Sequence-to-sequence models for chatbots 19 Utility landscape

Applications of Evolutionary Computation - Kevin Sim 2018-03-07 This book constitutes the refereed conference proceedings of the 21st International Conference on the Applications of Evolutionary Computation, EvoApplications 2018, held in Parma, Italy, in April 2018, collocated with the Evo* 2018 events EuroGP, EvoCOP, and EvoMUSART. The 59 revised full papers presented were carefully reviewed and selected from 84 submissions. EvoApplications 2018 combined research from 14 different domains: business analytics and finance (EvoBAFIN); computational biology (EvoBIO); communication networks and other parallel and distributed systems (EvoCOMNET); complex systems (EvoCOMPLEX); energy-related optimization (EvoENERGY); games and multi-agent systems (EvoGAMES); image analysis, signal processing and pattern recognition (EvoIASP); realworld industrial and commercial environments (EvoINDUSTRY); knowledge incorporation in evolutionary computation (EvoKNOW); continuous parameter optimization (EvoNUM); parallel architectures and distributed infrastructures (EvoPAR); evolutionary robotics (EvoROBOT); nature-inspired algorithms in software engineering and testing (EvoSET); and stochastic and dynamic environments (EvoSTOC).

[Optoelectronics in Machine Vision-Based Theories and Applications](#) -

Rivas-Lopez, Moises 2018-08-17

Sensor technologies play a large part in modern life, as they are present in things like security systems, digital cameras, smartphones, and motion sensors. While these devices are always evolving, research is being done to further develop this technology to help detect and analyze threats, perform in-depth inspections, and perform tracking services.

Optoelectronics in Machine Vision-Based Theories and Applications provides innovative insights on theories and applications of optoelectronics in machine vision-based systems. It also covers topics such as applications of unmanned aerial vehicle, autonomous and mobile robots, medical scanning, industrial applications, agriculture, and structural health monitoring. This publication is a vital reference source for engineers, technology developers, academicians, researchers, and advanced-level students seeking emerging research on sensor technologies and machine vision.

The Book of Alternative Data - Alexander Denev 2020-06-29

The first and only book to systematically address methodologies and processes of leveraging non-traditional information sources in the context of investing and risk management Harnessing non-traditional data sources to generate alpha, analyze markets, and forecast risk is a subject of intense interest for financial professionals. A growing number of regularly-held conferences on alternative data are being established, complemented by an upsurge in new papers on the subject. Alternative data is starting to be steadily incorporated by conventional institutional investors and risk managers throughout the financial world.

Methodologies to analyze and extract value from alternative data, guidance on how to source data and integrate data flows within existing systems is currently not treated in literature. Filling this significant gap in knowledge, *The Book of Alternative Data* is the first and only book to offer a coherent, systematic treatment of the subject. This groundbreaking volume provides readers with a roadmap for navigating the complexities of an array of alternative data sources, and delivers the appropriate techniques to analyze them. The authors—leading experts in financial modeling, machine learning, and quantitative research and analytics—employ a step-by-step approach to guide readers through the dense jungle of generated data. A first-of-its kind treatment of alternative data types, sources, and methodologies, this innovative book: Provides an integrated modeling approach to extract value from multiple types of datasets Treats the processes needed to make alternative data signals operational Helps investors and risk managers rethink how they engage with alternative datasets Features practical use case studies in many different financial markets and real-world techniques Describes how to avoid potential pitfalls and missteps in starting the alternative data journey Explains how to integrate information from different datasets to maximize informational value *The Book of Alternative Data* is an indispensable resource for anyone wishing to analyze or monetize different non-traditional datasets, including Chief Investment Officers, Chief Risk Officers, risk professionals, investment professionals, traders, economists, and machine learning developers and users.

Biomedical and Business Applications Using Artificial Neural Networks and Machine Learning - Segall, Richard S. 2022-01-07

During these uncertain and turbulent times, intelligent technologies including artificial neural networks (ANN) and machine learning (ML) have played an incredible role in being able to predict, analyze, and navigate unprecedented circumstances across a number of industries, ranging from healthcare to hospitality. Multi-factor prediction in particular has been especially helpful in dealing with the most current pressing issues such as COVID-19 prediction, pneumonia detection, cardiovascular diagnosis and disease management, automobile accident prediction, and vacation rental listing analysis. To date, there has not been much research content readily available in these areas, especially content written extensively from a user perspective. *Biomedical and Business Applications Using Artificial Neural Networks and Machine Learning* is designed to cover a brief and focused range of essential topics in the field with perspectives, models, and first-hand experiences shared by prominent researchers, discussing applications of artificial neural networks (ANN) and machine learning (ML) for biomedical and business applications and a listing of current open-source software for neural networks, machine learning, and artificial intelligence. It also presents summaries of currently available open source software that utilize neural networks and machine learning. The book is ideal for professionals, researchers, students, and practitioners who want to more fully understand in a brief and concise format the realm and technologies of artificial neural networks (ANN) and machine learning (ML) and how they have been used for prediction of multi-disciplinary research

problems in a multitude of disciplines.

Information Systems and Management Science - Lalit Garg 2022-12-30

This multidisciplinary book delves into information systems' concepts, principles, methods and procedures and their innovative applications in management science and other domains, including business, industry, health care and education. It will be valuable to students, researchers, academicians, developers, policymakers and managers thriving to improve their information and management systems, develop new strategies to solve complex problems and implement novel techniques to utilise the massive data best. This book of Information Systems and Management Science (proceedings of ISMS 2021) is intended to be used as a reference by scholars, scientists and practitioners who collect scientific and technical contributions concerning models, tools, technologies and applications in the field of information systems and management science. This book shows how to exploit information systems in a technology-rich management field.

Deep Learning for Unmanned Systems - Anis Koubaa 2021-10-01

This book is used at the graduate or advanced undergraduate level and many others. Manned and unmanned ground, aerial and marine vehicles enable many promising and revolutionary civilian and military applications that will change our life in the near future. These applications include, but are not limited to, surveillance, search and rescue, environment monitoring, infrastructure monitoring, self-driving cars, contactless last-mile delivery vehicles, autonomous ships, precision agriculture and transmission line inspection to name just a few. These vehicles will benefit from advances of deep learning as a subfield of machine learning able to endow these vehicles with different capability such as perception, situation awareness, planning and intelligent control. Deep learning models also have the ability to generate actionable insights into the complex structures of large data sets. In recent years, deep learning research has received an increasing amount of attention from researchers in academia, government laboratories and industry. These research activities have borne some fruit in tackling some of the challenging problems of manned and unmanned ground, aerial and marine vehicles that are still open. Moreover, deep learning methods have been recently actively developed in other areas of machine learning, including reinforcement training and transfer/meta-learning, whereas standard, deep learning methods such as recent neural network (RNN) and coevolutionary neural networks (CNN). The book is primarily meant for researchers from academia and industry, who are working on in the research areas such as engineering, control engineering, robotics, mechatronics, biomedical engineering, mechanical engineering and computer science. The book chapters deal with the recent research problems in the areas of reinforcement learning-based control of UAVs and deep learning for unmanned aerial systems (UAS) The book chapters present various techniques of deep learning for robotic applications. The book chapters contain a good literature survey with a long list of references. The book chapters are well written with a good exposition of the research problem, methodology, block diagrams and mathematical techniques. The book chapters are lucidly illustrated with numerical examples and simulations. The book chapters discuss details of applications and future research areas.

Deep Learning for Beginners - Steven Cooper 2018-11-06

☐☐The Best Deep Learning Book for Beginners☐☐ If you are looking for a complete beginners guide to learn deep learning with examples, in just a few hours, then you need to continue reading. This book delves into the basics of deep learning for those who are enthusiasts concerning all things machine learning and artificial intelligence. For those who have seen movies which show computer systems taking over the world like, Terminator, or benevolent systems that watch over the population, i.e. Person of Interest, this should be right up your alley. This book will give you the basics of what deep learning entails. That means frameworks used by coders and significant components and tools used in deep learning, that enable facial recognition, speech recognition, and virtual assistance. Yes, deep learning provides the tools through which systems like Siri became possible. ☐☐ Grab your copy today and learn ☐☐ ♦ Deep learning utilizes frameworks which allow people to develop tools which are able to offer better abstraction, along with simplification of hard programming issues. TensorFlow is the most popular tool and is used by corporate giants such as Airbus, Twitter, and even Google. ♦ The book illustrates TensorFlow and Caffe2 as the prime frameworks that are used for development by Google and Facebook. Facebook illustrates Caffe2 as one of the lightweight and modular deep learning frameworks, though TensorFlow is the most popular one, considering it has a lot of popularity, and thus, a big forum, which allows for assistance on main

problems. ♦ The book considers several components and tools of deep learning such as the neural networks; CNNs, RNNs, GANs, and auto-encoders. These algorithms create the building blocks which propel deep learning and advance it. ♦ The book also considers several applications, including chatbots and virtual assistants, which have become the main focus for deep learning into the future, as they represent the next frontier in information gathering and connectivity. The Internet of Things is also represented here, as deep learning allows for the integration of various systems via an artificial intelligence system, which is already being used for the home and car functions. ♦ And much more... The use of data science adds a lot of value to businesses, and we will continue to see the need for data scientists grow. This book is probably one of the best books for beginners. It's a step-by-step guide for any person who wants to start learning deep learning and artificial intelligence from scratch. When data science can reduce spending costs by billions of dollars in the healthcare industry, why wait to jump in? If you want to get started on deep learning and the concepts that run artificial technologies, don't wait any longer. Scroll up and click the buy now button to get this book today!

HCI International 2020 - Late Breaking Papers: Multimodality and Intelligence - Constantine Stephanidis 2020-10-16

This book constitutes late breaking papers from the 22nd International Conference on Human-Computer Interaction, HCII 2020, which was held in July 2020. The conference was planned to take place in Copenhagen, Denmark, but had to change to a virtual conference mode due to the COVID-19 pandemic. From a total of 6326 submissions, a total of 1439 papers and 238 posters have been accepted for publication in the HCII 2020 proceedings before the conference took place. In addition, a total of 333 papers and 144 posters are included in the volumes of the proceedings published after the conference as "Late Breaking Work" (papers and posters). These contributions address the latest research and development efforts in the field and highlight the human aspects of design and use of computing systems.

Hands-On Deep Learning with TensorFlow - Dan Van Boxel 2017-07-31

This book is your guide to exploring the possibilities in the field of deep learning, making use of Google's TensorFlow. You will learn about convolutional neural networks, and logistic regression while training models for deep learning to gain key insights into your data. About This Book Explore various possibilities with deep learning and gain amazing insights from data using Google's brainchild-- TensorFlow Want to learn what more can be done with deep learning? Explore various neural networks with the help of this comprehensive guide Rich in concepts, advanced guide on deep learning that will give you background to innovate in your environment Who This Book Is For If you are a data scientist who performs machine learning on a regular basis, are familiar with deep neural networks, and now want to gain expertise in working with convoluted neural networks, then this book is for you. Some familiarity with C++ or Python is assumed. What You Will Learn Set up your computing environment and install TensorFlow Build simple TensorFlow graphs for everyday computations Apply logistic regression for classification with TensorFlow Design and train a multilayer neural network with TensorFlow Intuitively understand convolutional neural networks for image recognition Bootstrap a neural network from simple to more accurate models See how to use TensorFlow with other types of networks Program networks with SciKit-Flow, a high-level interface to TensorFlow In Detail Dan Van Boxel's Deep Learning with TensorFlow is based on Dan's best-selling TensorFlow video course. With deep learning going mainstream, making sense of data and getting accurate results using deep networks is possible. Dan Van Boxel will be your guide to exploring the possibilities with deep learning; he will enable you to understand data like never before. With the efficiency and simplicity of TensorFlow, you will be able to process your data and gain insights that will change how you look at data. With Dan's guidance, you will dig deeper into the hidden layers of abstraction using raw data. Dan then shows you various complex algorithms for deep learning and various examples that use these deep neural networks. You will also learn how to train your machine to craft new features to make sense of deeper layers of data. In this book, Dan shares his knowledge across topics such as logistic regression, convolutional neural networks, recurrent neural networks, training deep networks, and high level interfaces. With the help of novel practical examples, you will become an ace at advanced multilayer networks, image recognition, and beyond. Style and Approach This book is your go-to guide to becoming a deep learning expert in your organization. Dan helps you evaluate common and not-so-common deep

neural networks with the help of insightful examples that you can relate to, and show how they can be exploited in the real world with complex raw data.

The Democratization of Artificial Intelligence - Andreas Sudmann 2019-10-31

After a long time of neglect, Artificial Intelligence is once again at the center of most of our political, economic, and socio-cultural debates. Recent advances in the field of Artificial Neural Networks have led to a renaissance of dystopian and utopian speculations on an AI-rendered future. Algorithmic technologies are deployed for identifying potential terrorists through vast surveillance networks, for producing sentencing guidelines and recidivism risk profiles in criminal justice systems, for demographic and psychographic targeting of bodies for advertising or propaganda, and more generally for automating the analysis of language, text, and images. Against this background, the aim of this book is to discuss the heterogeneous conditions, implications, and effects of modern AI and Internet technologies in terms of their political dimension: What does it mean to critically investigate efforts of net politics in the age of machine learning algorithms?

TensorFlow for Deep Learning - Bharath Ramsundar 2018-03-01

Learn how to solve challenging machine learning problems with TensorFlow, Google's revolutionary new software library for deep learning. If you have some background in basic linear algebra and calculus, this practical book introduces machine-learning fundamentals by showing you how to design systems capable of detecting objects in images, understanding text, analyzing video, and predicting the properties of potential medicines. TensorFlow for Deep Learning teaches concepts through practical examples and helps you build knowledge of deep learning foundations from the ground up. It's ideal for practicing developers with experience designing software systems, and useful for scientists and other professionals familiar with scripting but not necessarily with designing learning algorithms. Learn TensorFlow fundamentals, including how to perform basic computation Build simple learning systems to understand their mathematical foundations Dive into fully connected deep networks used in thousands of applications Turn prototypes into high-quality models with hyperparameter optimization Process images with convolutional neural networks Handle natural language datasets with recurrent neural networks Use reinforcement learning to solve games such as tic-tac-toe Train deep networks with hardware including GPUs and tensor processing units

Mastering TensorFlow 1. X - Armando Fandango 2018-01-22

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, TF Slim, 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.