

Getting Started Tensorflow Giancarlo Zaccone

A comprehensive guide to help you understand the principles of Reactive and asynchronous programming and its benefits

Key Features

- Explore the advantages of Reactive programming
- Use concurrency and parallelism in RxPY to build powerful reactive applications
- Deploy and scale your reactive applications using Docker

Book Description

Reactive programming is central to many concurrent systems, but it's famous for its steep learning curve, which makes most developers feel like they're hitting a wall. With this book, you will get to grips with reactive programming by steadily exploring various concepts

This hands-on guide gets you started with Reactive Programming (RP) in Python. You will learn about the principles and benefits of using RP, which can be leveraged to build powerful concurrent applications. As you progress through the chapters, you will be introduced to the paradigm of Functional and Reactive Programming (FaRP), observables and observers, and concurrency and parallelism.

The book will then take you through the implementation of an audio transcoding server and introduce you to a library that helps in the writing of FaRP code. You will understand how to use third-party services and dynamically reconfigure an application. By the end of the book, you will also have learned how to deploy and scale your applications with Docker and Traefik and explore the significant potential behind the reactive streams concept, and you'll have got to grips with a comprehensive set of best practices.

What you will learn

- Structure Python code for better readability, testing, and performance
- Explore the world of event-based programming
- Grasp the use of the most common operators in Rx
- Understand reactive extensions beyond simple examples
- Master the art of writing reusable components
- Deploy an application on a cloud platform with Docker and Traefik

Who this book is for

If you are a Python developer who wants to learn Reactive programming to build powerful concurrent and asynchronous applications, this book is for you. Basic understanding of the Python language is all you need to understand the concepts covered in this book.

Explore TensorFlow's capabilities to perform efficient deep learning on images

Key Features

- Discover image processing for machine vision
- Build an effective image classification system using the power of CNNs
- Leverage TensorFlow's capabilities to perform efficient deep learning

Book Description

TensorFlow is Google's popular offering for machine learning and deep learning, quickly becoming a favorite tool for performing fast, efficient, and accurate deep learning tasks. Hands-On Deep Learning for Images with TensorFlow shows you the practical implementations of real-world projects, teaching you how to leverage TensorFlow's capabilities to perform efficient image processing using the power of deep learning. With the help of this book, you will get to grips with the different paradigms of performing deep learning such as deep neural nets and convolutional neural networks, followed by understanding how they can be implemented using TensorFlow. By the end of this book, you will have mastered all the concepts of deep learning and their implementation with TensorFlow and Keras.

What you will learn

- Build machine learning models particularly focused on the MNIST digits
- Work with Docker and Keras to build an image classifier
- Understand natural language models to process text and images
- Prepare your dataset for machine learning
- Create classical, convolutional, and deep neural networks
- Create a RESTful image classification server

Who this book is for

Hands-On Deep Learning for Images with

TensorFlow is for you if you are an application developer, data scientist, or machine learning practitioner looking to integrate machine learning into application software and master deep learning by implementing practical projects in TensorFlow. Knowledge of Python programming and basics of deep learning are required to get the best out of this book.

Learn how to use Python to create efficient applications About This Book Identify the bottlenecks in your applications and solve them using the best profiling techniques Write efficient numerical code in NumPy, Cython, and Pandas Adapt your programs to run on multiple processors and machines with parallel programming Who This Book Is For The book is aimed at Python developers who want to improve the performance of their application. Basic knowledge of Python is expected What You Will Learn Write efficient numerical code with the NumPy and Pandas libraries Use Cython and Numba to achieve native performance Find bottlenecks in your Python code using profilers Write asynchronous code using Asyncio and RxPy Use Tensorflow and Theano for automatic parallelism in Python Set up and run distributed algorithms on a cluster using Dask and PySpark In Detail Python is a versatile language that has found applications in many industries. The clean syntax, rich standard library, and vast selection of third-party libraries make Python a wildly popular language. Python High Performance is a practical guide that shows how to leverage the power of both native and third-party Python libraries to build robust applications. The book explains how to use various profilers to find performance bottlenecks and apply the correct algorithm to fix them. The reader will learn how to effectively use NumPy and Cython to speed up numerical code. The book explains concepts of concurrent programming and how to implement robust and responsive applications using Reactive programming. Readers will learn how to write code for parallel architectures using Tensorflow and Theano, and use a cluster of computers for large-scale computations using technologies such as Dask and PySpark. By the end of the book, readers will have learned to achieve performance and scale from their Python applications. Style and approach A step-by-step practical guide filled with real-world use cases and examples

Learn how to use TensorFlow 2.0 to build machine learning and deep learning models with complete examples. The book begins with introducing TensorFlow 2.0 framework and the major changes from its last release. Next, it focuses on building Supervised Machine Learning models using TensorFlow 2.0. It also demonstrates how to build models using customer estimators. Further, it explains how to use TensorFlow 2.0 API to build machine learning and deep learning models for image classification using the standard as well as custom parameters. You'll review sequence predictions, saving, serving, deploying, and standardized datasets, and then deploy these models to production. All the code presented in the book will be available in the form of executable scripts at Github which allows you to try out the examples and extend them in interesting ways. What You'll Learn Review the new features of TensorFlow 2.0 Use TensorFlow 2.0 to build machine learning and deep learning models Perform sequence predictions using TensorFlow 2.0 Deploy TensorFlow 2.0 models with practical examples Who This Book Is For Data scientists, machine and deep learning engineers. Neural Networks and their implementation decoded with TensorFlow About This Book Develop a strong background in neural network programming from scratch, using the popular Tensorflow library. Use Tensorflow to implement different kinds of neural

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

Build machine and deep learning systems with the newly released TensorFlow 2 and Keras for the lab, production, and mobile devices Key Features Introduces and then uses TensorFlow 2 and Keras right from the start Teaches key machine and deep learning techniques Understand the fundamentals of deep learning and machine learning through clear explanations and extensive code samples Book Description Deep Learning with TensorFlow 2 and Keras, Second Edition teaches neural networks and deep learning techniques alongside TensorFlow (TF) and Keras. You'll learn how

to write deep learning applications in the most powerful, popular, and scalable machine learning stack available. TensorFlow is the machine learning library of choice for professional applications, while Keras offers a simple and powerful Python API for accessing TensorFlow. TensorFlow 2 provides full Keras integration, making advanced machine learning easier and more convenient than ever before. This book also introduces neural networks with TensorFlow, runs through the main applications (regression, ConvNets (CNNs), GANs, RNNs, NLP), covers two working example apps, and then dives into TF in production, TF mobile, and using TensorFlow with AutoML. What you will learn Build machine learning and deep learning systems with TensorFlow 2 and the Keras API Use Regression analysis, the most popular approach to machine learning Understand ConvNets (convolutional neural networks) and how they are essential for deep learning systems such as image classifiers Use GANs (generative adversarial networks) to create new data that fits with existing patterns Discover RNNs (recurrent neural networks) that can process sequences of input intelligently, using one part of a sequence to correctly interpret another Apply deep learning to natural human language and interpret natural language texts to produce an appropriate response Train your models on the cloud and put TF to work in real environments Explore how Google tools can automate simple ML workflows without the need for complex modeling Who this book is for This book is for Python developers and data scientists who want to build machine learning and deep learning systems with TensorFlow. Whether or not you have done machine learning before, this book gives you the theory and practice required to use Keras, TensorFlow 2, and AutoML to build machine learning systems. This practical guide will teach you how deep learning (DL) can be used to solve complex real-world problems. Key Features Explore deep reinforcement learning (RL), from the first principles to the latest algorithms Evaluate high-profile RL methods, including value iteration, deep Q-networks, policy gradients, TRPO, PPO, DDPG, D4PG, evolution strategies and genetic algorithms Keep up with the very latest industry developments, including AI-driven chatbots Book Description Recent developments in reinforcement learning (RL), combined with deep learning (DL), have seen unprecedented progress made towards training agents to solve complex problems in a human-like way. Google's use of algorithms to play and defeat the well-known Atari arcade games has propelled the field to prominence, and researchers are generating new ideas at a rapid pace. Deep Reinforcement Learning Hands-On is a comprehensive guide to the very latest DL tools and their limitations. You will evaluate methods including Cross-entropy and policy gradients, before applying them to real-world environments. Take on both the Atari set of virtual games and family favorites such as Connect4. The book provides an introduction to the basics of RL, giving you the know-how to code intelligent learning agents to take on a formidable array of practical tasks. Discover how to implement Q-learning on 'grid world' environments, teach your agent to buy and trade stocks, and find out how natural language models are driving the boom in chatbots. What you will learn Understand the DL context of RL and implement complex DL models Learn the foundation of RL: Markov decision processes Evaluate RL methods including Cross-entropy, DQN, Actor-Critic, TRPO, PPO, DDPG, D4PG and others Discover how to deal with discrete and continuous action spaces in various environments Defeat Atari arcade games using the value iteration method Create your own OpenAI Gym environment to train a stock trading

agent Teach your agent to play Connect4 using AlphaGo Zero Explore the very latest deep RL research on topics including AI-driven chatbots Who this book is for Some fluency in Python is assumed. Basic deep learning (DL) approaches should be familiar to readers and some practical experience in DL will be helpful. This book is an introduction to deep reinforcement learning (RL) and requires no background in RL. Summary Machine Learning with TensorFlow gives readers a solid foundation in machine-learning concepts plus hands-on experience coding TensorFlow with Python. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology TensorFlow, Google's library for large-scale machine learning, simplifies often-complex computations by representing them as graphs and efficiently mapping parts of the graphs to machines in a cluster or to the processors of a single machine. About the Book Machine Learning with TensorFlow gives readers a solid foundation in machine-learning concepts plus hands-on experience coding TensorFlow with Python. You'll learn the basics by working with classic prediction, classification, and clustering algorithms. Then, you'll move on to the money chapters: exploration of deep-learning concepts like autoencoders, recurrent neural networks, and reinforcement learning. Digest this book and you will be ready to use TensorFlow for machine-learning and deep-learning applications of your own. What's Inside Matching your tasks to the right machine-learning and deep-learning approaches Visualizing algorithms with TensorBoard Understanding and using neural networks About the Reader Written for developers experienced with Python and algebraic concepts like vectors and matrices. About the Author Author Nishant Shukla is a computer vision researcher focused on applying machine-learning techniques in robotics. Senior technical editor, Kenneth Fricklas, is a seasoned developer, author, and machine-learning practitioner. Table of Contents PART 1 - YOUR MACHINE-LEARNING RIG A machine-learning odyssey TensorFlow essentials PART 2 - CORE LEARNING ALGORITHMS Linear regression and beyond A gentle introduction to classification Automatically clustering data Hidden Markov models PART 3 - THE NEURAL NETWORK PARADIGM A peek into autoencoders Reinforcement learning Convolutional neural networks Recurrent neural networks Sequence-to-sequence models for chatbots Utility landscape Delve into neural networks, implement deep learning algorithms, and explore layers of data abstraction with the help of TensorFlow. 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 Gain real-world contextualization through some deep learning problems concerning research and application Who This Book Is For The book is for 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 Apply deep machine intelligence and GPU computing with TensorFlow Access public datasets and use TensorFlow to load, process, and transform the data Discover how to use the high-level TensorFlow API to build more powerful applications Use 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 a branch of machine learning algorithms based on learning multiple levels of abstraction. Neural networks, which are at the core of deep learning, are being used in predictive analytics, computer vision, natural language processing, time series forecasting, and to perform a myriad of other complex tasks. This book is conceived for developers, data analysts, machine learning practitioners and deep learning enthusiasts who want to build powerful, robust, and accurate predictive models with the power of TensorFlow, combined with other open source Python libraries. Throughout the book, you'll learn how to develop deep learning applications for machine learning systems using Feedforward Neural Networks, Convolutional Neural Networks, Recurrent Neural Networks, Autoencoders, and Factorization Machines. Discover how to attain deep learning programming on GPU in a distributed way. You'll come away with an in-depth knowledge of machine learning techniques and the skills to apply them to real-world projects. Style and approach This step-by-step guide explores common, and not so common, deep neural networks, and shows ho ...

Skip the theory and get the most out of Tensorflow to build production-ready machine learning models Key Features Exploit the features of Tensorflow to build and deploy machine learning models Train neural networks to tackle real-world problems in Computer Vision and NLP Handy techniques to write production-ready code for your Tensorflow models Book Description 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 allow you to dig deeper and gain more insights into your data than ever before. With the help of this book, you will work with recipes for training models, model evaluation, sentiment analysis, regression analysis, clustering analysis, artificial neural networks, and more. You will explore RNNs, CNNs, GANs, reinforcement learning, and capsule networks, each using Google's machine learning library, TensorFlow. Through real-world examples, you will get hands-on experience with linear regression techniques with TensorFlow. Once you are familiar and comfortable with the TensorFlow ecosystem, you will be shown how to take it to production. By the end of the book, you will be proficient in the field of machine intelligence using TensorFlow. You will also have good insight into deep learning and be capable of implementing machine learning algorithms in real-world scenarios. What you will learn Become familiar with the basic features of the TensorFlow library Get to know Linear Regression techniques with TensorFlow Learn SVMs with hands-on recipes Implement neural networks to improve predictive modeling Apply NLP and sentiment analysis to your data Master CNN and RNN through practical recipes Implement the gradient boosted random forest to predict housing prices Take TensorFlow into production Who this book is for If you are a data scientist or a machine learning engineer with some knowledge of linear algebra, statistics, and machine learning, this book is for you. If you want to skip the theory and build production-ready machine learning models using Tensorflow without reading pages and pages of material, this book is for you. Some background in Python programming is assumed.

Learn advanced analytical techniques and leverage existing tool kits to make your analytic applications more powerful, precise, and efficient. This book provides the right combination of architecture, design, and implementation information to create analytical systems that go beyond the basics of classification, clustering, and recommendation. Pro Hadoop Data Analytics emphasizes best practices to ensure coherent, efficient development. A complete example system will be developed using standard third-party components that consist of the tool kits, libraries, visualization and reporting code, as well as support glue to provide a working and extensible end-to-end system. The book also highlights the importance of end-to-end,

flexible, configurable, high-performance data pipeline systems with analytical components as well as appropriate visualization results. You'll discover the importance of mix-and-match or hybrid systems, using different analytical components in one application. This hybrid approach will be prominent in the examples. What You'll Learn Build big data analytic systems with the Hadoop ecosystem Use libraries, tool kits, and algorithms to make development easier and more effective Apply metrics to measure performance and efficiency of components and systems Connect to standard relational databases, noSQL data sources, and more Follow case studies with example components to create your own systems Who This Book Is For Software engineers, architects, and data scientists with an interest in the design and implementation of big data analytical systems using Hadoop, the Hadoop ecosystem, and other associated technologies.

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

quality of results with transfer learning Build TensorFlow models that can scale to large datasets and systems Who this book is for This book is for Software Engineers, Data Scientists, or Machine Learning practitioners who want to use CNNs for solving real-world problems. Knowledge of basic machine learning concepts, linear algebra and Python will help. A short and straight to the point guide that explains the implementation of Regular Expressions in Python. This book is aimed at Python developers who want to learn how to leverage Regular Expressions in Python. Basic knowledge of Python is required for a better understanding. Transform Your Ideas into High-Quality Python Code! Zed Shaw has perfected the world's best system for becoming a truly effective Python 3.x developer. Follow it and you will succeed—just like the tens of millions of programmers he's already taught. You bring the discipline, commitment, and persistence; the author supplies everything else. In Learn Python 3 the Hard Way, Zed Shaw taught you the basics of Programming with Python 3. Now, in Learn More Python 3 the Hard Way, you'll go far beyond the basics by working through 52 brilliantly crafted projects. Each one helps you build a key practical skill, combining demos to get you started and challenges to deepen your understanding. Zed then teaches you even more in 12 hours of online videos, where he shows you how to break, fix, and debug your code. First, you'll discover how to analyze a concept, idea, or problem to implement in software. Then, step by step, you'll learn to design solutions based on your analyses and implement them as simply and elegantly as possible. Throughout, Shaw stresses process so you can get started and build momentum, creativity to solve new problems, and quality so you'll build code people can rely on. Manage complex projects with a programmer's text editor Leverage the immense power of data structures Apply algorithms to process your data structures Master indispensable text parsing and processing techniques Use SQL to efficiently and logically model stored data Learn powerful command-line tools and skills Combine multiple practices in complete projects It'll be hard at first. But soon, you'll just get it—and that will feel great! This course will reward you for every minute you put into it. Soon, you'll go beyond merely writing code that runs: you'll craft high-quality Python code that solves real problems. You'll be a serious Python programmer. Perfect for Everyone Who's Already Started Working with Python, including Junior Developers and Seasoned Python Programmers Upgrading to Python 3.6+ Register your product at informit.com/register for convenient access to downloads, updates, and/or corrections as they become available.

Master efficient parallel programming to build powerful applications using Python About This Book Design and implement efficient parallel software Master new programming techniques to address and solve complex programming problems Explore the world of parallel programming with this book, which is a go-to resource for different kinds of parallel computing tasks in Python, using examples and topics covered in great depth Who This Book Is For Python Parallel Programming Cookbook is intended for software developers who are well versed with Python and want to use parallel programming techniques to write powerful and efficient code. This book will help you master the basics and the advanced of parallel computing. What You Will Learn Synchronize multiple threads and processes to manage parallel tasks Implement message passing communication between processes to build parallel applications Program your own GPU cards to address complex problems Manage computing entities to execute distributed computational tasks Write efficient programs by adopting the event-driven programming model Explore the cloud technology with Django and Google App Engine Apply parallel programming techniques that can lead to performance improvements In Detail Parallel programming techniques are required for a developer to get the best use of all the computational resources available today and to build efficient software systems. From multi-core to GPU systems up to the distributed architectures, the high computation of programs throughout requires the use of programming tools and software libraries. Because of this, it is becoming increasingly important to know what the parallel programming techniques are.

Python is commonly used as even non-experts can easily deal with its concepts. This book will teach you parallel programming techniques using examples in Python and will help you explore the many ways in which you can write code that allows more than one process to happen at once. Starting with introducing you to the world of parallel computing, it moves on to cover the fundamentals in Python. This is followed by exploring the thread-based parallelism model using the Python threading module by synchronizing threads and using locks, mutex, semaphores queues, GIL, and the thread pool. Next you will be taught about process-based parallelism where you will synchronize processes using message passing along with learning about the performance of MPI Python Modules. You will then go on to learn the asynchronous parallel programming model using the Python asyncio module along with handling exceptions. Moving on, you will discover distributed computing with Python, and learn how to install a broker, use Celery Python Module, and create a worker. You will also understand the StarCluster framework, Pycsp, Scoop, and Disco modules in Python. Further on, you will learn GPU programming with Python using the PyCUDA module along with evaluating performance limitations. Next you will get acquainted with the cloud computing concepts in Python, using Google App Engine (GAE), and building your first application with GAE. Lastly, you will learn about grid computing concepts in Python and using PyGlobus toolkit, GFTP and GASS COPY to transfer files, and service monitoring in PyGlobus. Style and approach A step-by-step guide to parallel programming using Python, with recipes accompanied by one or more programming examples. It is a practically oriented book and has all the necessary underlying parallel computing concepts.

A comprehensive guide to advanced deep learning techniques, including Autoencoders, GANs, VAEs, and Deep Reinforcement Learning, that drive today's most impressive AI results

Key Features Explore the most advanced deep learning techniques that drive modern AI results Implement Deep Neural Networks, Autoencoders, GANs, VAEs, and Deep Reinforcement Learning A wide study of GANs, including Improved GANs, Cross-Domain GANs and Disentangled Representation GANs

Book Description Recent developments in deep learning, including GANs, Variational Autoencoders, and Deep Reinforcement Learning, are creating impressive AI results in our news headlines - such as AlphaGo Zero beating world chess champions, and generative AI that can create art paintings that sell for over \$400k because they are so human-like. Advanced Deep Learning with Keras is a comprehensive guide to the advanced deep learning techniques available today, so you can create your own cutting-edge AI. Using Keras as an open-source deep learning library, you'll find hands-on projects throughout that show you how to create more effective AI with the latest techniques. The journey begins with an overview of MLPs, CNNs, and RNNs, which are the building blocks for the more advanced techniques in the book. You'll learn how to implement deep learning models with Keras and Tensorflow, and move forwards to advanced techniques, as you explore deep neural network architectures, including ResNet and DenseNet, and how to create Autoencoders. You then learn all about Generative Adversarial Networks (GANs), and how they can open new levels of AI performance. Variational AutoEncoders (VAEs) are implemented, and you'll see how GANs and VAEs have the generative power to synthesize data that can be extremely convincing to humans - a major stride forward for modern AI. To complete this set of advanced techniques, you'll learn how to implement Deep Reinforcement Learning (DRL) such as Deep Q-Learning and Policy Gradient Methods, which are critical to many modern results in AI. What you will learn

Cutting-edge techniques in human-like AI performance Implement advanced deep learning models using Keras The building blocks for advanced techniques - MLPs, CNNs, and RNNs Deep neural networks – ResNet and DenseNet Autoencoders and Variational AutoEncoders (VAEs) Generative Adversarial Networks (GANs) and creative AI techniques Disentangled Representation GANs, and Cross-Domain GANs Deep Reinforcement Learning (DRL) methods and implementation Produce

industry-standard applications using OpenAI gym Deep Q-Learning and Policy Gradient Methods Who this book is for Some fluency with Python is assumed. As an advanced book, you'll be familiar with some machine learning approaches, and some practical experience with DL will be helpful. Knowledge of Keras or TensorFlow is not required but would be helpful.

Deep Learning with TensorFlow Packt Publishing Ltd

Perform advanced data manipulation tasks using pandas and become an expert data analyst. Key Features Manipulate and analyze your data expertly using the power of pandas Work with missing data and time series data and become a true pandas expert Includes expert tips and techniques on making your data analysis tasks easier Book Description pandas is a popular Python library used by data scientists and analysts worldwide to manipulate and analyze their data. This book presents useful data manipulation techniques in pandas to perform complex data analysis in various domains. An update to our highly successful previous edition with new features, examples, updated code, and more, this book is an in-depth guide to get the most out of pandas for data analysis. Designed for both intermediate users as well as seasoned practitioners, you will learn advanced data manipulation techniques, such as multi-indexing, modifying data structures, and sampling your data, which allow for powerful analysis and help you gain accurate insights from it. With the help of this book, you will apply pandas to different domains, such as Bayesian statistics, predictive analytics, and time series analysis using an example-based approach. And not just that; you will also learn how to prepare powerful, interactive business reports in pandas using the Jupyter notebook. By the end of this book, you will learn how to perform efficient data analysis using pandas on complex data, and become an expert data analyst or data scientist in the process. What you will learn Speed up your data analysis by importing data into pandas Keep relevant data points by selecting subsets of your data Create a high-quality dataset by cleaning data and fixing missing values Compute actionable analytics with grouping and aggregation in pandas Master time series data analysis in pandas Make powerful reports in pandas using Jupyter notebooks Who this book is for This book is for data scientists, analysts and Python developers who wish to explore advanced data analysis and scientific computing techniques using pandas. Some fundamental understanding of Python programming and familiarity with the basic data analysis concepts is all you need to get started with this book.

The purpose of this book is two-fold, we focus on detailed coverage of deep learning and transfer learning, comparing and contrasting the two with easy-to-follow concepts and examples. The second area of focus is on real-world examples and research problems using TensorFlow, Keras and Python ecosystem with hands-on examples.

Step-by-step guide to learn and solve complex computational problems with Nature Inspired algorithms. DESCRIPTION Natural Computing is the field of research inspired by nature, that allows the development of new algorithms to solve complex problems, leads to the synthesis of natural models, and may result in the design of new computing systems. This book exactly aims to educate you with practical examples on topics of importance associated with research field of Natural computing. The initial few chapters will quickly walk you through Neural Networks while describing deep learning architectures such as CNN, RNN and AutoEncoders using Keras. As you progress further, you'll gain understanding to develop genetic algorithm to solve traveling salesman problem, implement swarm intelligence techniques using the SwarmPackagePy and Cellular Automata techniques such as Game of Life, Langton's ant, etc. The latter half of the book will introduce you to the world of Fractals such as the Cantor Set and the Mandelbrot Set, develop a quantum program with the QiSkit tool that runs on a real quantum computing platform, namely the IBM Q Machine and a Python simulation of the Adleman experiment that showed for the first time the possibility of performing computations at the molecular level. KEY FEATURES Artificial Neural Networks Deep Learning models using Keras Quantum Computers and Programming Genetic Algorithms,

CNN and RNNs Swarm Intelligence Systems Reinforcement Learning using OpenAI Artificial Life DNA computing Fractals WHAT WILL YOU LEARN Mastering Artificial Neural Networks Developing Artificial Intelligence systems Resolving complex problems with Genetic Programming and Swarm intelligence algorithms Programming Quantum Computers Exploring the mathematical world of fractals Simulating complex systems by Cellular Automata Understanding the basics of DNA computation WHO THIS BOOK IS FOR This book is for all science enthusiasts, in particular who want to understand what are the links between computer sciences and natural systems. Interested readers should have good skills in math and python programming along with some basic knowledge of physics and biology. . Although, some knowledge of the topics covered in the book will be helpful, it is not essential to have worked with the tools covered in the book. Table of Contents Neural Networks Deep Learning Genetic Programming Swarm Intelligence Cellular Automata Fractals Quantum Computing DNA Computing

Leverage the power of the Reinforcement Learning techniques to develop self-learning systems using Tensorflow Key Features Learn reinforcement learning concepts and their implementation using TensorFlow Discover different problem-solving methods for Reinforcement Learning Apply reinforcement learning for autonomous driving cars, roboadvisors, and more Book Description Reinforcement Learning (RL), allows you to develop smart, quick and self-learning systems in your business surroundings. It is an effective method to train your learning agents and solve a variety of problems in Artificial Intelligence—from games, self-driving cars and robots to enterprise applications that range from datacenter energy saving (cooling data centers) to smart warehousing solutions. The book covers the major advancements and successes achieved in deep reinforcement learning by synergizing deep neural network architectures with reinforcement learning. The book also introduces readers to the concept of Reinforcement Learning, its advantages and why it's gaining so much popularity. The book also discusses on MDPs, Monte Carlo tree searches, dynamic programming such as policy and value iteration, temporal difference learning such as Q-learning and SARSA. You will use TensorFlow and OpenAI Gym to build simple neural network models that learn from their own actions. You will also see how reinforcement learning algorithms play a role in games, image processing and NLP. By the end of this book, you will have a firm understanding of what reinforcement learning is and how to put your knowledge to practical use by leveraging the power of TensorFlow and OpenAI Gym. What you will learn Implement state-of-the-art Reinforcement Learning algorithms from the basics Discover various techniques of Reinforcement Learning such as MDP, Q Learning and more Learn the applications of Reinforcement Learning in advertisement, image processing, and NLP Teach a Reinforcement Learning model to play a game using TensorFlow and the OpenAI gym Understand how Reinforcement Learning Applications are used in robotics Who this book is for If you want to get started with reinforcement learning using TensorFlow in the most practical way, this book will be a useful resource. The book assumes prior knowledge of machine learning and neural network programming concepts, as well as some understanding of the TensorFlow framework. No previous experience with Reinforcement Learning is required. Implement effective programming techniques in Python to build scalable software that saves time and memory Key Features Design distributed computing systems and massive computational tasks coherently Learn practical recipes with concise explanations that address development pain points encountered while coding parallel programs Understand how to host your parallelized applications on the cloud Book Description Nowadays, it has become extremely important for programmers to understand the link between the software and the parallel nature of their hardware so that their programs run efficiently on computer architectures. Applications based on parallel programming are fast, robust, and easily scalable. This updated edition features cutting-edge techniques for building effective concurrent

applications in Python 3.7. The book introduces parallel programming architectures and covers the fundamental recipes for thread-based and process-based parallelism. You'll learn about mutex, semaphores, locks, queues exploiting the threading, and multiprocessing modules, all of which are basic tools to build parallel applications. Recipes on MPI programming will help you to synchronize processes using the fundamental message passing techniques with `mpi4py`. Furthermore, you'll get to grips with asynchronous programming and how to use the power of the GPU with PyCUDA and PyOpenCL frameworks. Finally, you'll explore how to design distributed computing systems with Celery and architect Python apps on the cloud using PythonAnywhere, Docker, and serverless applications. By the end of this book, you will be confident in building concurrent and high-performing applications in Python. What you will learn

- Synchronize multiple threads and processes to manage parallel tasks
- Use message passing techniques to establish communication between processes to build parallel applications
- Program your own GPU cards to address complex problems
- Manage computing entities to execute distributed computational task
- Write efficient programs by adopting the event-driven programming model
- Explore cloud technology with Django and Google App Engine
- Apply parallel programming techniques that can lead to performance improvements

Who this book is for The Python Parallel Programming Cookbook is for software developers who are well-versed with Python and want to use parallel programming techniques to write powerful and efficient code. This book will help you master the basics and the advanced of parallel computing.

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.

Delve into neural networks, implement deep learning algorithms, and explore layers of data abstraction with the help of TensorFlow. Key Features 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 Gain real-world contextualization through some deep learning problems concerning research and application Book Description Deep learning is a branch of machine learning algorithms based on learning multiple levels of abstraction. Neural networks, which are at the core of deep learning, are being used in predictive analytics, computer vision, natural language processing, time series forecasting, and to perform a myriad of other complex tasks. This book is conceived for developers, data analysts, machine learning practitioners and deep learning enthusiasts who want to build powerful, robust, and accurate predictive models with the power of TensorFlow, combined with other open source Python libraries. Throughout the book, you'll learn how to develop deep learning applications for machine learning systems using Feedforward Neural Networks, Convolutional Neural Networks, Recurrent Neural Networks, Autoencoders, and Factorization Machines. Discover how to attain deep learning programming on GPU in a distributed way. You'll come away with an in-depth knowledge of machine learning techniques and the skills to apply them to real-world projects. What you will learn Apply deep machine intelligence and GPU computing with TensorFlow Access public datasets and use TensorFlow to load, process, and transform the data Discover how to use the high-level TensorFlow API to build more powerful applications Use 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 Who this book is for The book is for 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.

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 deep learning, Natural Language Processing (NLP), speech recognition, and general predictive analytics. The book provides a hands-on approach to TensorFlow fundamentals for a broad technical audience—from data scientists and engineers to students and researchers. The authors begin by working through some basic examples in TensorFlow before diving deeper into topics such as CNN, RNN, LSTM, and GNN. The book is written for those who want to build powerful, robust, and accurate predictive models with the power of TensorFlow, combined with other open source Python libraries. The authors demonstrate TensorFlow projects on Single Board Computers (SBCs).

Create distributed applications with clever design patterns to solve complex problems Key Features Set up and run distributed algorithms on a cluster using Dask and PySpark Master skills to accurately implement concurrency in your code Gain practical experience of Python design patterns with real-world examples Book Description This Learning Path shows you how to leverage the power of both native and third-party Python libraries for building robust and responsive applications. You will learn about profilers and reactive programming,

concurrency and parallelism, as well as tools for making your apps quick and efficient. You will discover how to write code for parallel architectures using TensorFlow and Theano, and use a cluster of computers for large-scale computations using technologies such as Dask and PySpark. With the knowledge of how Python design patterns work, you will be able to clone objects, secure interfaces, dynamically choose algorithms, and accomplish much more in high performance computing. By the end of this Learning Path, you will have the skills and confidence to build engaging models that quickly offer efficient solutions to your problems. This Learning Path includes content from the following Packt products: Python High Performance - Second Edition by Gabriele Lanaro Mastering Concurrency in Python by Quan Nguyen Mastering Python Design Patterns by Sakis Kasampalis What you will learn Use NumPy and pandas to import and manipulate datasets Achieve native performance with Cython and Numba Write asynchronous code using asyncio and RxPy Design highly scalable programs with application scaffolding Explore abstract methods to maintain data consistency Clone objects using the prototype pattern Use the adapter pattern to make incompatible interfaces compatible Employ the strategy pattern to dynamically choose an algorithm Who this book is for This Learning Path is specially designed for Python developers who want to build high-performance applications and learn about single core and multi-core programming, distributed concurrency, and Python design patterns. Some experience with Python programming language will help you get the most out of this Learning Path.

Intended to anyone interested in numerical computing and data science: students, researchers, teachers, engineers, analysts, hobbyists... Basic knowledge of Python/NumPy is recommended. Some skills in mathematics will help you understand the theory behind the computational methods.

Get command of your organizational Big Data using the power of data science and analytics Key Features A perfect companion to boost your Big Data storing, processing, analyzing skills to help you take informed business decisions Work with the best tools such as Apache Hadoop, R, Python, and Spark for NoSQL platforms to perform massive online analyses Get expert tips on statistical inference, machine learning, mathematical modeling, and data visualization for Big Data Book Description Big Data analytics relates to the strategies used by organizations to collect, organize and analyze large amounts of data to uncover valuable business insights that otherwise cannot be analyzed through traditional systems. Crafting an enterprise-scale cost-efficient Big Data and machine learning solution to uncover insights and value from your organization's data is a challenge. Today, with hundreds of new Big Data systems, machine learning packages and BI Tools, selecting the right combination of technologies is an even greater challenge. This book will help you do that. With the help of this guide, you will be able to bridge the gap between the theoretical world of technology with the practical ground reality of building corporate Big Data and

data science platforms. You will get hands-on exposure to Hadoop and Spark, build machine learning dashboards using R and R Shiny, create web-based apps using NoSQL databases such as MongoDB and even learn how to write R code for neural networks. By the end of the book, you will have a very clear and concrete understanding of what Big Data analytics means, how it drives revenues for organizations, and how you can develop your own Big Data analytics solution using different tools and methods articulated in this book. What you will learn - Get a 360-degree view into the world of Big Data, data science and machine learning - Broad range of technical and business Big Data analytics topics that caters to the interests of the technical experts as well as corporate IT executives - Get hands-on experience with industry-standard Big Data and machine learning tools such as Hadoop, Spark, MongoDB, KDB+ and R - Create production-grade machine learning BI Dashboards using R and R Shiny with step-by-step instructions - Learn how to combine open-source Big Data, machine learning and BI Tools to create low-cost business analytics applications - Understand corporate strategies for successful Big Data and data science projects - Go beyond general-purpose analytics to develop cutting-edge Big Data applications using emerging technologies Who this book is for The book is intended for existing and aspiring Big Data professionals who wish to become the go-to person in their organization when it comes to Big Data architecture, analytics, and governance. While no prior knowledge of Big Data or related technologies is assumed, it will be helpful to have some programming experience.

Updated and revised second edition of the bestselling guide to advanced deep learning with TensorFlow 2 and Keras Key Features Explore the most advanced deep learning techniques that drive modern AI results New coverage of unsupervised deep learning using mutual information, object detection, and semantic segmentation Completely updated for TensorFlow 2.x Book Description Advanced Deep Learning with TensorFlow 2 and Keras, Second Edition is a completely updated edition of the bestselling guide to the advanced deep learning techniques available today. Revised for TensorFlow 2.x, this edition introduces you to the practical side of deep learning with new chapters on unsupervised learning using mutual information, object detection (SSD), and semantic segmentation (FCN and PSPNet), further allowing you to create your own cutting-edge AI projects. Using Keras as an open-source deep learning library, the book features hands-on projects that show you how to create more effective AI with the most up-to-date techniques. Starting with an overview of multi-layer perceptrons (MLPs), convolutional neural networks (CNNs), and recurrent neural networks (RNNs), the book then introduces more cutting-edge techniques as you explore deep neural network architectures, including ResNet and DenseNet, and how to create autoencoders. You will then learn about GANs, and how they can unlock new levels of AI performance. Next, you'll discover how a variational autoencoder (VAE) is implemented, and how GANs and VAEs have the generative power to synthesize data that can be extremely convincing to

humans. You'll also learn to implement DRL such as Deep Q-Learning and Policy Gradient Methods, which are critical to many modern results in AI. What you will learn Use mutual information maximization techniques to perform unsupervised learning Use segmentation to identify the pixel-wise class of each object in an image Identify both the bounding box and class of objects in an image using object detection Learn the building blocks for advanced techniques - MLPs, CNN, and RNNs Understand deep neural networks - including ResNet and DenseNet Understand and build autoregressive models – autoencoders, VAEs, and GANs Discover and implement deep reinforcement learning methods Who this book is for This is not an introductory book, so fluency with Python is required. The reader should also be familiar with some machine learning approaches, and practical experience with DL will also be helpful. Knowledge of Keras or TensorFlow 2.0 is not required but is recommended.

This book is a collection of notes and sample codes written by the author while he was learning Neural Networks in Machine Learning. Topics include Neural Networks (NN) concepts: nodes, layers, activation functions, learning rates, training sets, etc.; deep playground for classical neural networks; building neural networks with Python; walking through Tariq Rashi's 'Make Your Own Neural Network' source code; using 'TensorFlow' and 'PyTorch' machine learning platforms; understanding CNN (Convolutional Neural Network), RNN (Recurrent Neural Network), GNN (Graph Neural Network). Updated in 2020 (Version 1.20) with Deep Playground tutorials. For latest updates and free sample chapters, visit <http://www.herongyang.com/Neural-Network>.

Simplify next-generation deep learning by implementing powerful generative models using Python, TensorFlow and Keras Key Features Understand the common architecture of different types of GANs Train, optimize, and deploy GAN applications using TensorFlow and Keras Build generative models with real-world data sets, including 2D and 3D data Book Description Developing Generative Adversarial Networks (GANs) is a complex task, and it is often hard to find code that is easy to understand. This book leads you through eight different examples of modern GAN implementations, including CycleGAN, simGAN, DCGAN, and 2D image to 3D model generation. Each chapter contains useful recipes to build on a common architecture in Python, TensorFlow and Keras to explore increasingly difficult GAN architectures in an easy-to-read format. The book starts by covering the different types of GAN architecture to help you understand how the model works. This book also contains intuitive recipes to help you work with use cases involving DCGAN, Pix2Pix, and so on. To understand these complex applications, you will take different real-world data sets and put them to use. By the end of this book, you will be equipped to deal with the challenges and issues that you may face while working with GAN models, thanks to easy-to-follow code solutions that you can implement right away. What you will learn Structure a GAN architecture in pseudocode Understand the common architecture for each of the GAN models you will build Implement different GAN architectures in TensorFlow

and Keras Use different datasets to enable neural network functionality in GAN models Combine different GAN models and learn how to fine-tune them Produce a model that can take 2D images and produce 3D models Develop a GAN to do style transfer with Pix2Pix Who this book is for This book is for data scientists, machine learning developers, and deep learning practitioners looking for a quick reference to tackle challenges and tasks in the GAN domain. Familiarity with machine learning concepts and working knowledge of Python programming language will help you get the most out of the book.

Write modern natural language processing applications using deep learning algorithms and TensorFlow Key Features Focuses on more efficient natural language processing using TensorFlow Covers NLP as a field in its own right to improve understanding for choosing TensorFlow tools and other deep learning approaches Provides choices for how to process and evaluate large unstructured text datasets Learn to apply the TensorFlow toolbox to specific tasks in the most interesting field in artificial intelligence Book Description Natural language processing (NLP) supplies the majority of data available to deep learning applications, while TensorFlow is the most important deep learning framework currently available. Natural Language Processing with TensorFlow brings TensorFlow and NLP together to give you invaluable tools to work with the immense volume of unstructured data in today's data streams, and apply these tools to specific NLP tasks. Thushan Ganegedara starts by giving you a grounding in NLP and TensorFlow basics. You'll then learn how to use Word2vec, including advanced extensions, to create word embeddings that turn sequences of words into vectors accessible to deep learning algorithms. Chapters on classical deep learning algorithms, like convolutional neural networks (CNN) and recurrent neural networks (RNN), demonstrate important NLP tasks as sentence classification and language generation. You will learn how to apply high-performance RNN models, like long short-term memory (LSTM) cells, to NLP tasks. You will also explore neural machine translation and implement a neural machine translator. After reading this book, you will gain an understanding of NLP and you'll have the skills to apply TensorFlow in deep learning NLP applications, and how to perform specific NLP tasks. What you will learn Core concepts of NLP and various approaches to natural language processing How to solve NLP tasks by applying TensorFlow functions to create neural networks Strategies to process large amounts of data into word representations that can be used by deep learning applications Techniques for performing sentence classification and language generation using CNNs and RNNs About employing state-of-the art advanced RNNs, like long short-term memory, to solve complex text generation tasks How to write automatic translation programs and implement an actual neural machine translator from scratch The trends and innovations that are paving the future in NLP Who this book is for This book is for Python developers with a strong interest in deep learning, who want to learn how to leverage TensorFlow to simplify NLP tasks.

Fundamental Python skills are assumed, as well as some knowledge of machine learning and undergraduate-level calculus and linear algebra. No previous natural language processing experience required, although some background in NLP or computational linguistics will be helpful.

Explore GPU-enabled programmable environment for machine learning, scientific applications, and gaming using PuCUDA, PyOpenGL, and Anaconda Accelerate
Key Features Understand effective synchronization strategies for faster processing using GPUs Write parallel processing scripts with PyCuda and PyOpenCL Learn to use the CUDA libraries like CuDNN for deep learning on GPUs
Book Description GPUs are proving to be excellent general purpose-parallel computing solutions for high performance tasks such as deep learning and scientific computing. This book will be your guide to getting started with GPU computing. It will start with introducing GPU computing and explain the architecture and programming models for GPUs. You will learn, by example, how to perform GPU programming with Python, and you'll look at using integrations such as PyCUDA, PyOpenCL, CuPy and Numba with Anaconda for various tasks such as machine learning and data mining. Going further, you will get to grips with GPU work flows, management, and deployment using modern containerization solutions. Toward the end of the book, you will get familiar with the principles of distributed computing for training machine learning models and enhancing efficiency and performance. By the end of this book, you will be able to set up a GPU ecosystem for running complex applications and data models that demand great processing capabilities, and be able to efficiently manage memory to compute your application effectively and quickly. What you will learn Utilize Python libraries and frameworks for GPU acceleration Set up a GPU-enabled programmable machine learning environment on your system with Anaconda Deploy your machine learning system on cloud containers with illustrated examples Explore PyCUDA and PyOpenCL and compare them with platforms such as CUDA, OpenCL and ROCm. Perform data mining tasks with machine learning models on GPUs Extend your knowledge of GPU computing in scientific applications Who this book is for Data Scientist, Machine Learning enthusiasts and professionals who wants to get started with GPU computation and perform the complex tasks with low-latency. Intermediate knowledge of Python programming is assumed.

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. Grasp the fundamental concepts of deep learning using Tensorflow in a hands-on manner
Key Features Get a first-hand experience of the deep learning concepts and techniques with this easy-to-follow guide Train different types of neural networks using Tensorflow for real-world problems in language

processing, computer vision, transfer learning, and more. Designed for those who believe in the concept of 'learn by doing', this book is a perfect blend of theory and code examples. **Book Description** Deep learning is a popular subset of machine learning, and it allows you to build complex models that are faster and give more accurate predictions. This book is your companion to take your first steps into the world of deep learning, with hands-on examples to boost your understanding of the topic. This book starts with a quick overview of the essential concepts of data science and machine learning which are required to get started with deep learning. It introduces you to Tensorflow, the most widely used machine learning library for training deep learning models. You will then work on your first deep learning problem by training a deep feed-forward neural network for digit classification, and move on to tackle other real-world problems in computer vision, language processing, sentiment analysis, and more. Advanced deep learning models such as generative adversarial networks and their applications are also covered in this book. By the end of this book, you will have a solid understanding of all the essential concepts in deep learning. With the help of the examples and code provided in this book, you will be equipped to train your own deep learning models with more confidence. **What you will learn** Understand the fundamentals of deep learning and how it is different from machine learning Get familiarized with Tensorflow, one of the most popular libraries for advanced machine learning Increase the predictive power of your model using feature engineering Understand the basics of deep learning by solving a digit classification problem of MNIST Demonstrate face generation based on the CelebA database, a promising application of generative models Apply deep learning to other domains like language modeling, sentiment analysis, and machine translation **Who this book is for** This book targets data scientists and machine learning developers who wish to get started with deep learning. If you know what deep learning is but are not quite sure of how to use it, this book will help you as well. An understanding of statistics and data science concepts is required. Some familiarity with Python programming will also be beneficial.

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

Get up and running with the latest numerical computing library by Google and dive deeper into your data! About This Book- Get the first book on the market that shows you the key aspects TensorFlow, how it works, and how to use it for the second generation of machine learning- Want to perform faster and more accurate computations in the field of data science? This book will acquaint you with an all-new refreshing library-TensorFlow!- Dive into the next generation of numerical computing and get the most out of your data with this quick guide Who This Book Is For This book is dedicated to all the machine learning and deep learning enthusiasts, data scientists, researchers, and even students who want to perform more accurate, fast machine learning operations with TensorFlow. Those with basic knowledge of programming (Python and C/C++) and math concepts who want to be introduced to the topics of machine learning will find this book useful. What You Will Learn- Install and adopt TensorFlow in your Python environment to solve mathematical problems- Get to know the basic machine and deep learning concepts- Train and test neural networks to fit your data model- Make predictions using regression algorithms- Analyze your data with a clustering procedure- Develop algorithms for clustering and data classification- Use GPU computing to analyze big data In Detail Google's TensorFlow engine, after much fanfare, has evolved in to a robust, user-friendly, and customizable, application-grade software library of machine learning (ML) code for numerical computation and neural networks. This book takes you through the practical software implementation of various machine learning techniques with TensorFlow. In the first few chapters, you'll gain familiarity with the framework and perform the mathematical operations required for data analysis. As you progress further, you'll learn to implement various machine learning techniques such as classification, clustering, neural networks, and deep learning through practical examples. By the end of this book, you'll have gained hands-on experience of using TensorFlow and building classification, image recognition systems, language processing, and information retrieving systems for your application. Style and approach Get quickly up and running with TensorFlow using this fast-paced guide. You will get to know everything that can be done with TensorFlow and we'll show you how to implement it in your environment. The examples in the book are from the core of the computation industry-something you can connect to and will find familiar.

The purpose of this book is to help to spread TensorFlow knowledge among engineers who want to expand their wisdom in the exciting world of Machine

Learning. We believe that anyone with an engineering background might require from now on Deep Learning, an

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.

[Copyright: 2a2788f8530f42de3ca78ace520cdf69](#)