

## Pro Deep Learning With Tensorflow A Mathematical Approach To Advanced Artificial Intelligence In Python

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.

## Read Free Pro Deep Learning With Tensorflow A Mathematical Approach To Advanced Artificial Intelligence In Python

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.

NVIDIA's Full-Color Guide to Deep Learning: All Students Need to Get Started and Get Results Learning Deep Learning is a complete guide to DL. Illuminating both the core concepts and the hands-on programming techniques needed to succeed, this book suits seasoned developers, data scientists, analysts, but also those with no prior machine learning or statistic experience. After introducing the essential building blocks of deep neural networks, such as artificial neurons and fully connected, convolutional, and recurrent layers, Magnus Ekman shows how to use them to build advanced architectures, including the Transformer. He describes how these concepts are used to build modern networks for

## Read Free Pro Deep Learning With Tensorflow A Mathematical Approach To Advanced Artificial Intelligence In Python

computer vision and natural language processing (NLP), including Mask R-CNN, GPT, and BERT. And he explains how a natural language translator and a system generating natural language descriptions of images. Throughout, Ekman provides concise, well-annotated code examples using TensorFlow with Keras. Corresponding PyTorch examples are provided online, and the book thereby covers the two dominating Python libraries for DL used in industry and academia. He concludes with an introduction to neural architecture search (NAS), exploring important ethical issues and providing resources for further learning. Explore and master core concepts: perceptrons, gradient-based learning, sigmoid neurons, and back propagation See how DL frameworks make it easier to develop more complicated and useful neural networks Discover how convolutional neural networks (CNNs) revolutionize image classification and analysis Apply recurrent neural networks (RNNs) and long short-term memory (LSTM) to text and other variable-length sequences Master NLP with sequence-to-sequence networks and the Transformer architecture Build applications for natural language translation and image captioning

Take a hands-on approach to understanding deep learning and build smart applications that can recognize images and interpret text

**Key Features** Understand how to implement deep learning with TensorFlow and Keras Learn the fundamentals of computer vision and image recognition Study the architecture of different neural networks

**Book Description** Are you fascinated by how deep learning powers intelligent applications such as self-driving cars, virtual assistants, facial recognition devices, and chatbots to process data and solve complex problems? Whether you are familiar with machine learning or are new to this domain, The Deep Learning Workshop will make it easy for you to understand deep learning with the help of interesting examples and exercises throughout. The book starts by highlighting the relationship between deep learning, machine learning, and artificial intelligence and helps you get comfortable with the TensorFlow 2.0 programming structure using hands-on exercises. You'll understand neural networks, the structure of a perceptron, and how to use TensorFlow to create and train models. The book will then let you explore the fundamentals of computer vision by performing image recognition exercises with convolutional neural networks (CNNs) using Keras. As you advance, you'll be able to make your model more powerful by implementing text embedding and sequencing the data using popular deep learning solutions. Finally, you'll get to grips with bidirectional recurrent neural networks (RNNs) and build generative adversarial networks (GANs) for image synthesis. By the end of this deep learning book, you'll have learned the skills essential for building deep learning models with TensorFlow and Keras. What you will learn

Understand how deep learning, machine learning, and artificial intelligence are different Develop multilayer deep neural networks with TensorFlow Implement deep neural networks for multiclass classification using Keras Train CNN models for image recognition Handle sequence data and use it in conjunction with RNNs Build a GAN to generate high-quality synthesized

## Read Free Pro Deep Learning With Tensorflow A Mathematical Approach To Advanced Artificial Intelligence In Python

images Who this book is for If you are interested in machine learning and want to create and train deep learning models using TensorFlow and Keras, this workshop is for you. A solid understanding of Python and its packages, along with basic machine learning concepts, will help you to learn the topics quickly.

Understand the principles and practices of machine learning and deep learning This hands-on guide lays out machine learning and deep learning techniques and technologies in a style that is approachable, using just the basic math required. Written by a pair of experts in the field, Machine Learning and Deep Learning Using Python and TensorFlow contains case studies in several industries, including banking, insurance, e-commerce, retail, and healthcare. The book shows how to utilize machine learning and deep learning functions in today's smart devices and apps. You will get download links for datasets, code, and sample projects referred to in the text. Coverage includes: Machine learning and deep learning concepts Python programming and statistics fundamentals Regression and logistic regression Decision trees Model selection and cross-validation Cluster analysis Random forests and boosting Artificial neural networks TensorFlow and Keras Deep learning hyperparameters Convolutional neural networks Recurrent neural networks and long short-term memory

Develop machine learning models across various domains. This book offers a single source that provides comprehensive coverage of the capabilities of TensorFlow 2 through the use of realistic, scenario-based projects. After learning what's new in TensorFlow 2, you'll dive right into developing machine learning models through applicable projects. This book covers a wide variety of ANN architectures—starting from working with a simple sequential network to advanced CNN, RNN, LSTM, DCGAN, and so on. A full chapter is devoted to each kind of network and each chapter consists of a full project describing the network architecture used, the theory behind that architecture, what data set is used, the pre-processing of data, model training, testing and performance optimizations, and analysis. This practical approach can either be used from the beginning through to the end or, if you're already familiar with basic ML models, you can dive right into the application that interests you. Line-by-line explanations on major code segments help to fill in the details as you work and the entire project source is available to you online for learning and further experimentation. With Artificial Neural Networks with TensorFlow 2 you'll see just how wide the range of TensorFlow's capabilities are. What You'll Learn Develop Machine Learning Applications Translate languages using neural networks Compose images with style transfer Who This Book Is For Beginners, practitioners, and hard-cored developers who want to master machine and deep learning with TensorFlow 2. The reader should have working concepts of ML basics and terminologies.

This book shows readers how they can successfully analyze data using only two core machine learning algorithms---and how to do so using the popular Python programming language. These algorithms deal with common scenarios faced by

## Read Free Pro Deep Learning With Tensorflow A Mathematical Approach To Advanced Artificial Intelligence In Python

all data analysts and data scientists. This book focuses on two algorithm families (linear methods and ensemble methods) that effectively predict outcomes. This type of problem covers a multitude of use cases (what ad to place on a web page, predicting prices in securities markets, detecting credit card fraud, etc.). The focus on two families gives enough room for full descriptions of the mechanisms at work in the algorithms. Then the code examples serve to illustrate the workings of the machinery with specific hackable code. The author will explain in simple terms, using no complex math, how these algorithms work, and will then show how to apply them in Python. He will also provide advice on how to select from among these algorithms, and will show how to prepare the data, and how to use the trained models in practice. The author begins with an overview of the two core algorithms, explaining the types of problems solved by each one. He then introduces a core set of Python programming techniques that can be used to apply these algorithms. The author shows various techniques for building predictive models that solve a range of problems, from simple to complex; he also shows how to measure the performance of each model to ensure you use the right one. The following chapters provide a deep dive into each of the two algorithms: penalized linear regression and ensemble methods. Chapters will show how to apply each algorithm in Python. Readers can directly use the sample code to build their own solutions.

### Pro Deep Learning with TensorFlow A Mathematical Approach to Advanced Artificial Intelligence in Python

Whether you're a software engineer aspiring to enter the world of deep learning, a veteran data scientist, or a hobbyist with a simple dream of making the next viral AI app, you might have wondered where to begin. This step-by-step guide teaches you how to build practical deep learning applications for the cloud, mobile, browsers, and edge devices using a hands-on approach. Relying on years of industry experience transforming deep learning research into award-winning applications, Anirudh Koul, Siddha Ganju, and Meher Kasam guide you through the process of converting an idea into something that people in the real world can use. Train, tune, and deploy computer vision models with Keras, TensorFlow, Core ML, and TensorFlow Lite Develop AI for a range of devices including Raspberry Pi, Jetson Nano, and Google Coral Explore fun projects, from Silicon Valley's Not Hotdog app to 40+ industry case studies Simulate an autonomous car in a video game environment and build a miniature version with reinforcement learning Use transfer learning to train models in minutes Discover 50+ practical tips for maximizing model accuracy and speed, debugging, and scaling to millions of users

Discover more insight about deep learning and how to work with Swift for TensorFlow to develop intelligent apps. TensorFlow was designed for easy adoption by iOS programmers working in Swift. This book covers the established and tested concepts and ties them to modern Swift programming and applicable use in developing for iOS. Using illustrative examples, the book starts off by introducing you to basic machine learning concepts along with code snippets in Swift for TensorFlow.. Fundamentals of neural networks required to understand today's deep learning research will be covered and put in the context of working in the Swift language with the goal of developing primarily for Apple's mobile ecosystem. Other important topics covered include computation graphs, loss functions, optimization techniques, regularizing neural networks, recurrent neural networks—such as those used in Siri and Google Translate; and convolutional neural networks. You'll also learn to reuse pre-trained neural networks and work with generative models. Finally, developing and building in security to models is addressed. Swift

## Read Free Pro Deep Learning With Tensorflow A Mathematical Approach To Advanced Artificial Intelligence In Python

code will be provided throughout the book to keep the concepts grounded in application within Apple's frameworks. What You'll Learn • Write machine learning code in Swift • Run neural networks in Apple environments • Apply fundamental deep learning concepts to mobile app development Who This Book Is For Programmers familiar with Swift and the basics of AI

Artificial intelligence and its various components are rapidly engulfing almost every professional industry. Specific features of AI that have proven to be vital solutions to numerous real-world issues are machine learning and deep learning. These intelligent agents unlock higher levels of performance and efficiency, creating a wide span of industrial applications. However, there is a lack of research on the specific uses of machine/deep learning in the professional realm. Machine Learning and Deep Learning in Real-Time Applications provides emerging research exploring the theoretical and practical aspects of machine learning and deep learning and their implementations as well as their ability to solve real-world problems within several professional disciplines including healthcare, business, and computer science. Featuring coverage on a broad range of topics such as image processing, medical improvements, and smart grids, this book is ideally designed for researchers, academicians, scientists, industry experts, scholars, IT professionals, engineers, and students seeking current research on the multifaceted uses and implementations of machine learning and deep learning across the globe.

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.

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

# Read Free Pro Deep Learning With Tensorflow A Mathematical Approach To Advanced Artificial Intelligence In Python

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

This book is a step-by-step guide to show you how to implement generative models in TensorFlow 2.x from scratch. You'll get to grips with the image generative technology by covering autoencoders, style transfer, and GANs as well as fundamental and state-of-the-art models. Become a machine learning pro! Google TensorFlow has become the darling of financial firms and research organizations, but the technology can be intimidating and the learning curve is steep. Luckily, TensorFlow For Dummies is here to offer you a friendly, easy-to-follow book on the subject. Inside, you'll find out how to write applications with TensorFlow, while also grasping the concepts underlying machine learning—all without ever losing your cool! Machine learning has become ubiquitous in modern society, and its applications include language translation, robotics, handwriting analysis, financial prediction, and image recognition. TensorFlow is Google's preeminent toolset for machine learning, and this hands-on guide makes it easy to understand, even for those without a background in artificial intelligence. Install TensorFlow on your computer Learn the fundamentals of statistical regression and neural networks Visualize the machine learning process with TensorBoard Perform image recognition with convolutional neural networks (CNNs) Analyze sequential data with recurrent neural networks (RNNs) Execute TensorFlow on mobile devices and the Google Cloud Platform (GCP) If you're a manager or software developer looking to use TensorFlow for machine learning, this is the book you'll want to have close by.

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

## Read Free Pro Deep Learning With Tensorflow A Mathematical Approach To Advanced Artificial Intelligence In Python

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.

A project-based guide to the basics of deep learning. This concise, project-driven guide to deep learning takes readers through a series of program-writing tasks that introduce them to the use of deep learning in such areas of artificial intelligence as computer vision, natural-language processing, and reinforcement learning. The author, a longtime artificial intelligence researcher specializing in natural-language processing, covers feed-forward neural nets, convolutional neural nets, word embeddings, recurrent neural nets, sequence-to-sequence learning, deep reinforcement learning, unsupervised models, and other fundamental concepts and techniques. Students and practitioners learn the basics of deep learning by working through programs in Tensorflow, an open-source machine learning framework. "I find I learn computer science material best by sitting down and writing programs," the author writes, and the book reflects this approach. Each chapter includes a programming project, exercises, and references for further reading. An early chapter is devoted to Tensorflow and its interface with Python, the widely used programming language. Familiarity with linear algebra, multivariate calculus, and probability and statistics is required, as is a rudimentary knowledge of programming in Python. The book can be used in both undergraduate and graduate courses; practitioners will find it an essential reference.

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

## Read Free Pro Deep Learning With Tensorflow A Mathematical Approach To Advanced Artificial Intelligence In Python

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 book presents solutions to the majority of the challenges you will face while training neural networks to solve deep learning problems. It covers the trending deep learning architectures used in industry and tackles a variety of use cases in computer vision, text processing, audio analysis, recommender systems, and game bots 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

## Read Free Pro Deep Learning With Tensorflow A Mathematical Approach To Advanced Artificial Intelligence In Python

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

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 Shanging 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

Through a series of recent breakthroughs, deep learning has boosted the entire field of machine learning. Now, even

## Read Free Pro Deep Learning With Tensorflow A Mathematical Approach To Advanced Artificial Intelligence In Python

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. Explore the machine learning landscape, particularly neural nets Use Scikit-Learn to track an example machine-learning project end-to-end Explore several training models, including support vector machines, decision trees, random forests, and ensemble methods Use the TensorFlow library to build and train neural nets Dive into neural net architectures, including convolutional nets, recurrent nets, and deep reinforcement learning Learn techniques for training and scaling deep neural nets

Every other day we hear about new ways to put deep learning to good use: improved medical imaging, accurate credit card fraud detection, long range weather forecasting, and more. PyTorch puts these superpowers in your hands, providing a comfortable Python experience that gets you started quickly and then grows with you as you—and your deep learning skills—become more sophisticated. Deep Learning with PyTorch will make that journey engaging and fun. Summary Every other day we hear about new ways to put deep learning to good use: improved medical imaging, accurate credit card fraud detection, long range weather forecasting, and more. PyTorch puts these superpowers in your hands, providing a comfortable Python experience that gets you started quickly and then grows with you as you—and your deep learning skills—become more sophisticated. Deep Learning with PyTorch will make that journey engaging and fun. Foreword by Soumith Chintala, Cocreator of PyTorch. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the technology Although many deep learning tools use Python, the PyTorch library is truly Pythonic. Instantly familiar to anyone who knows PyData tools like NumPy and scikit-learn, PyTorch simplifies deep learning without sacrificing advanced features. It's excellent for building quick models, and it scales smoothly from laptop to enterprise. Because companies like Apple, Facebook, and JPMorgan Chase rely on PyTorch, it's a great skill to have as you expand your career options. It's easy to get started with PyTorch. It minimizes cognitive overhead without sacrificing the access to advanced features, meaning you can focus on what matters the most - building and training the latest and greatest deep learning models and contribute to making a dent in the world. PyTorch is also a snap to scale and extend, and it partners well with other Python tooling. PyTorch has been adopted by hundreds of deep learning practitioners and several first-class players like FAIR, OpenAI, FastAI and Purdue. About the book Deep Learning with PyTorch teaches you to create neural networks and deep learning systems with PyTorch. This practical

## Read Free Pro Deep Learning With Tensorflow A Mathematical Approach To Advanced Artificial Intelligence In Python

book quickly gets you to work building a real-world example from scratch: a tumor image classifier. Along the way, it covers best practices for the entire DL pipeline, including the PyTorch Tensor API, loading data in Python, monitoring training, and visualizing results. After covering the basics, the book will take you on a journey through larger projects. The centerpiece of the book is a neural network designed for cancer detection. You'll discover ways for training networks with limited inputs and start processing data to get some results. You'll sift through the unreliable initial results and focus on how to diagnose and fix the problems in your neural network. Finally, you'll look at ways to improve your results by training with augmented data, make improvements to the model architecture, and perform other fine tuning. What's inside

Training deep neural networks  
Implementing modules and loss functions  
Utilizing pretrained models from PyTorch Hub  
Exploring code samples in Jupyter Notebooks

About the reader  
For Python programmers with an interest in machine learning.

About the author  
Eli Stevens had roles from software engineer to CTO, and is currently working on machine learning in the self-driving-car industry. Luca Antiga is cofounder of an AI engineering company and an AI tech startup, as well as a former PyTorch contributor. Thomas Viehmann is a PyTorch core developer and machine learning trainer and consultant. consultant based in Munich, Germany and a PyTorch core developer.

Table of Contents

PART 1 - CORE PYTORCH

1 Introducing deep learning and the PyTorch Library

2 Pretrained networks

3 It starts with a tensor

4 Real-world data representation using tensors

5 The mechanics of learning

6 Using a neural network to fit the data

7 Telling birds from airplanes: Learning from images

8 Using convolutions to generalize

PART 2 - LEARNING FROM IMAGES IN THE REAL WORLD: EARLY DETECTION OF LUNG CANCER

9 Using PyTorch to fight cancer

10 Combining data sources into a unified dataset

11 Training a classification model to detect suspected tumors

12 Improving training with metrics and augmentation

13 Using segmentation to find suspected nodules

14 End-to-end nodule analysis, and where to go next

PART 3 - DEPLOYMENT

15 Deploying to production

Probabilistic Deep Learning is a hands-on guide to the principles that support neural networks. Learn to improve network performance with the right distribution for different data types, and discover Bayesian variants that can state their own uncertainty to increase accuracy. This book provides easy-to-apply code and uses popular frameworks to keep you focused on practical applications.

Summary Probabilistic Deep Learning: With Python, Keras and TensorFlow

Probability teaches the increasingly popular probabilistic approach to deep learning that allows you to refine your results more quickly and accurately without much trial-and-error testing. Emphasizing practical techniques that use the Python-based Tensorflow Probability Framework, you'll learn to build highly-performant deep learning applications that can reliably handle the noise and uncertainty of real-world data. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications.

About the technology  
The world is a noisy and uncertain place. Probabilistic

## Read Free Pro Deep Learning With Tensorflow A Mathematical Approach To Advanced Artificial Intelligence In Python

deep learning models capture that noise and uncertainty, pulling it into real-world scenarios. Crucial for self-driving cars and scientific testing, these techniques help deep learning engineers assess the accuracy of their results, spot errors, and improve their understanding of how algorithms work. About the book Probabilistic Deep Learning is a hands-on guide to the principles that support neural networks. Learn to improve network performance with the right distribution for different data types, and discover Bayesian variants that can state their own uncertainty to increase accuracy. This book provides easy-to-apply code and uses popular frameworks to keep you focused on practical applications. What's inside Explore maximum likelihood and the statistical basis of deep learning Discover probabilistic models that can indicate possible outcomes Learn to use normalizing flows for modeling and generating complex distributions Use Bayesian neural networks to access the uncertainty in the model About the reader For experienced machine learning developers. About the author Oliver Dürr is a professor at the University of Applied Sciences in Konstanz, Germany. Beate Sick holds a chair for applied statistics at ZHAW and works as a researcher and lecturer at the University of Zurich. Elvis Murina is a data scientist. Table of Contents PART 1 - BASICS OF DEEP LEARNING 1 Introduction to probabilistic deep learning 2 Neural network architectures 3 Principles of curve fitting PART 2 - MAXIMUM LIKELIHOOD APPROACHES FOR PROBABILISTIC DL MODELS 4 Building loss functions with the likelihood approach 5 Probabilistic deep learning models with TensorFlow Probability 6 Probabilistic deep learning models in the wild PART 3 - BAYESIAN APPROACHES FOR PROBABILISTIC DL MODELS 7 Bayesian learning 8 Bayesian neural networks

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

## Read Free Pro Deep Learning With Tensorflow A Mathematical Approach To Advanced Artificial Intelligence In Python

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 DetailDeep 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 approachThis 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.

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

## Read Free Pro Deep Learning With Tensorflow A Mathematical Approach To Advanced Artificial Intelligence In Python

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

Work through engaging and practical deep learning projects using TensorFlow 2.0. Using a hands-on approach, the projects in this book will lead new programmers through the basics into developing practical deep learning applications. Deep learning is quickly integrating itself into the technology landscape. Its applications range from applicable data science to deep fakes and so much more. It is crucial for aspiring data scientists or those who want to enter the field of AI to understand deep learning concepts. The best way to learn is by doing. You'll develop a working knowledge of not only TensorFlow, but also related technologies such as Python and Keras. You'll also work with Neural Networks and other deep learning concepts. By the end of the book, you'll have a collection of unique projects that you can add to your GitHub profiles and expand on for professional application. What You'll Learn Grasp the basic process of neural networks through projects, such as creating music Restore and colorize black and white images with deep learning processes Who This Book Is For Beginners new to TensorFlow and Python.

This book includes 9 projects on building smart and practical AI-based systems. These projects cover solutions to different domain-specific problems in healthcare, e-commerce and more. With this book, you will apply different machine learning and deep learning techniques and learn how to build your own intelligent applications for smart ...

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.

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

## Read Free Pro Deep Learning With Tensorflow A Mathematical Approach To Advanced Artificial Intelligence In Python

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

"The authors' clear visual style provides a comprehensive look at what's currently possible with artificial neural networks as well as a glimpse of the magic that's to come." —Tim Urban, author of *Wait But Why* Fully Practical, Insightful Guide to Modern Deep Learning Deep learning is transforming software, facilitating powerful new artificial intelligence capabilities, and driving unprecedented algorithm performance. *Deep Learning Illustrated* is uniquely intuitive and offers a complete introduction to the discipline's techniques. Packed with full-color figures and easy-to-follow code, it sweeps away the complexity of building deep learning models, making the subject approachable and fun to learn. World-class instructor and practitioner Jon Krohn—with visionary content from Grant Beyleveld and beautiful illustrations by Aglaé Bassens—presents straightforward analogies to explain what deep learning is, why it has become so popular, and how it relates to other machine learning approaches. Krohn has created a practical reference and tutorial for developers, data scientists, researchers, analysts, and students who want to start applying it. He illuminates theory with hands-on Python code in accompanying Jupyter notebooks. To help you progress quickly, he focuses on the versatile deep learning library Keras to nimbly construct efficient TensorFlow models; PyTorch, the leading alternative library, is also covered. You'll gain a pragmatic understanding of all major deep learning approaches and their uses in applications ranging from machine vision and natural language processing to image generation and game-playing algorithms. Discover what makes deep learning systems unique, and the implications for practitioners Explore new tools that make deep learning models easier to build, use, and improve Master essential theory: artificial neurons, training, optimization, convolutional nets, recurrent nets, generative adversarial networks (GANs), deep reinforcement learning, and more Walk through building interactive deep learning applications, and move forward with your own artificial intelligence projects Register your book for convenient access to downloads, updates, and/or corrections as they become available. See inside book

## Read Free Pro Deep Learning With Tensorflow A Mathematical Approach To Advanced Artificial Intelligence In Python

for details.

Summary Deep Learning with R introduces the world of deep learning using the powerful Keras library and its R language interface. The book builds your understanding of deep learning through intuitive explanations and practical examples. Continue your journey into the world of deep learning with Deep Learning with R in Motion, a practical, hands-on video course available exclusively at Manning.com ([www.manning.com/livevideo/deep-?learning-with-r-in-motion](http://www.manning.com/livevideo/deep-?learning-with-r-in-motion)). 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. Deep-learning systems now enable previously impossible smart applications, revolutionizing image recognition and natural-language processing, and identifying complex patterns in data. The Keras deep-learning library provides data scientists and developers working in R a state-of-the-art toolset for tackling deep-learning tasks. About the Book Deep Learning with R introduces the world of deep learning using the powerful Keras library and its R language interface. Initially written for Python as Deep Learning with Python by Keras creator and Google AI researcher François Chollet and adapted for R by RStudio founder J. J. Allaire, this book builds your understanding of deep learning through intuitive explanations and practical examples. You'll practice your new skills with R-based applications in computer vision, natural-language processing, and generative models. What's Inside Deep learning from first principles Setting up your own deep-learning environment Image classification and generation Deep learning for text and sequences About the Reader You'll need intermediate R programming skills. No previous experience with machine learning or deep learning is assumed. About the Authors François Chollet is a deep-learning researcher at Google and the author of the Keras library. J.J. Allaire is the founder of RStudio and the author of the R interfaces to TensorFlow and Keras. 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

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

## Read Free Pro Deep Learning With Tensorflow A Mathematical Approach To Advanced Artificial Intelligence In Python

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

Implement deep learning applications using TensorFlow while learning the “why” through in-depth conceptual explanations. You'll start by learning what deep learning offers over other machine learning models. Then familiarize yourself with several technologies used to create deep learning models. While some of these technologies are complementary, such as Pandas, Scikit-Learn, and Numpy—others are competitors, such as PyTorch, Caffe, and Theano. This book clarifies the positions of deep learning and Tensorflow among their peers. You'll then work on supervised deep learning models to gain applied experience with the technology. A single-layer of multiple perceptrons will be used to build a shallow neural network before turning it into a deep neural network. After showing the structure of the ANNs, a real-life application will be created with Tensorflow 2.0 Keras API. Next, you'll work on data augmentation and batch normalization methods. Then, the Fashion MNIST dataset will be used to train a CNN. CIFAR10 and Imagenet pre-

## Read Free Pro Deep Learning With Tensorflow A Mathematical Approach To Advanced Artificial Intelligence In Python

trained models will be loaded to create already advanced CNNs. Finally, move into theoretical applications and unsupervised learning with auto-encoders and reinforcement learning with tf-agent models. With this book, you'll delve into applied deep learning practical functions and build a wealth of knowledge about how to use TensorFlow effectively. What You'll Learn Compare competing technologies and see why TensorFlow is more popular Generate text, image, or sound with GANs Predict the rating or preference a user will give to an item Sequence data with recurrent neural networks Who This Book Is For Data scientists and programmers new to the fields of deep learning and machine learning APIs.

Deep learning networks are getting smaller. Much smaller. The Google Assistant team can detect words with a model just 14 kilobytes in size—small enough to run on a microcontroller. With this practical book you'll enter the field of TinyML, where deep learning and embedded systems combine to make astounding things possible with tiny devices. Pete Warden and Daniel Situnayake explain how you can train models small enough to fit into any environment. Ideal for software and hardware developers who want to build embedded systems using machine learning, this guide walks you through creating a series of TinyML projects, step-by-step. No machine learning or microcontroller experience is necessary. Build a speech recognizer, a camera that detects people, and a magic wand that responds to gestures Work with Arduino and ultra-low-power microcontrollers Learn the essentials of ML and how to train your own models Train models to understand audio, image, and accelerometer data Explore TensorFlow Lite for Microcontrollers, Google's toolkit for TinyML Debug applications and provide safeguards for privacy and security Optimize latency, energy usage, and model and binary size

This book is your guide to master deep learning with TensorFlow, with the help of 10 real-world projects. You will train high-performance models in TensorFlow to generate captions for images automatically, predict stocks' performance, create intelligent chatbots, perform large-scale text classification, develop recommendation systems, and more.

This friendly and accessible guide to AI theory and programming in Python requires no maths or data science background. Key Features Roll up your sleeves and start programming AI models No math, data science, or machine learning background required Packed with hands-on examples, illustrations, and clear step-by-step instructions 5 hands-on working projects put ideas into action and show step-by-step how to build intelligent software Book Description AI is changing the world – and with this book, anyone can start building intelligent software! Through his best-selling video courses, Hadelin de Ponteves has taught hundreds of thousands of people to write AI software. Now, for the first time, his hands-on, energetic approach is available as a book. Taking a graduated approach that starts with the basics before easing readers into more complicated formulas and notation, Hadelin helps you understand what you really need to build

## Read Free Pro Deep Learning With Tensorflow A Mathematical Approach To Advanced Artificial Intelligence In Python

AI systems with reinforcement learning and deep learning. Five full working projects put the ideas into action, showing step-by-step how to build intelligent software using the best and easiest tools for AI programming: Google Colab Python TensorFlow Keras PyTorch AI Crash Course teaches everyone to build an AI to work in their applications. Once you've read this book, you're only limited by your imagination. What you will learn Master the key skills of deep learning, reinforcement learning, and deep reinforcement learning Understand Q-learning and deep Q-learning Learn from friendly, plain English explanations and practical activities Build fun projects, including a virtual-self-driving car Use AI to solve real-world business problems and win classic video games Build an intelligent, virtual robot warehouse worker Who this book is for If you want to add AI to your skillset, this book is for you. It doesn't require data science or machine learning knowledge. Just maths basics (high school level).

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

# Read Free Pro Deep Learning With Tensorflow A Mathematical Approach To Advanced Artificial Intelligence In Python

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.

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

## Read Free Pro Deep Learning With Tensorflow A Mathematical Approach To Advanced Artificial Intelligence In Python

autoencoders Reinforcement learning Convolutional neural networks Recurrent neural networks Sequence-to-sequence models for chatbots Utility landscape

Implement TensorFlow's offerings such as TensorBoard, TensorFlow.js, TensorFlow Probability, and TensorFlow Lite to build smart automation projects Key Features Use machine learning and deep learning principles to build real-world projects Get to grips with TensorFlow's impressive range of module offerings Implement projects on GANs, reinforcement learning, and capsule network Book 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 learn Understand the TensorFlow ecosystem using various datasets and techniques Create recommendation systems for quality product recommendations Build projects using CNNs, NLP, and Bayesian neural networks Play Pac-Man using deep reinforcement learning Deploy scalable TensorFlow-based machine learning systems Generate your own book script using RNNs Who 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

[Copyright: 5112a15ebf1bde8b057d1b3d76191af8](#)