

Introduction To Tensorflow Python Programming Tutorials

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. This book gives you the theory and practice required to use Keras, TensorFlow 2, and AutoML to build machine learning systems. Some knowledge of machine learning is expected.*

Build real-world Artificial Intelligence applications with Python to intelligently interact with the world around you **About This Book** *Step into the amazing world of intelligent apps using this comprehensive guide* **Enter the world of Artificial Intelligence, explore it, and create your own applications** *Work through simple yet insightful examples that will get you up and running with Artificial Intelligence in no time* **Who This Book Is For** *This book is for Python developers who want to build real-world Artificial Intelligence applications. This book is friendly to Python beginners, but being familiar with Python would be useful to play around with the code. It will also be useful for experienced Python programmers who are looking to use Artificial Intelligence techniques in their existing technology stacks. What You Will Learn* **Realize different classification and regression techniques** *Understand the concept of clustering and how to use it to automatically segment data* **See how to build an intelligent recommender system**

Understand logic programming and how to use it Build automatic speech recognition systems Understand the basics of heuristic search and genetic programming Develop games using Artificial Intelligence Learn how reinforcement learning works Discover how to build intelligent applications centered on images, text, and time series data See how to use deep learning algorithms and build applications based on it In Detail Artificial Intelligence is becoming increasingly relevant in the modern world where everything is driven by technology and data. It is used extensively across many fields such as search engines, image recognition, robotics, finance, and so on. We will explore various real-world scenarios in this book and you'll learn about various algorithms that can be used to build Artificial Intelligence applications. During the course of this book, you will find out how to make informed decisions about what algorithms to use in a given context. Starting from the basics of Artificial Intelligence, you will learn how to develop various building blocks using different data mining techniques. You will see how to implement different algorithms to get the best possible results, and will understand how to apply them to real-world scenarios. If you want to add an intelligence layer to any application that's based on images, text, stock market, or some other form of data, this exciting book on Artificial Intelligence will definitely be your guide! Style and approach This highly practical book will show you how to implement Artificial Intelligence. The book provides multiple examples enabling you to create smart applications to meet the needs of your organization. In every chapter, we explain an algorithm, implement it, and then build a smart application.

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.

Master the world of Python and Machine Learning with this incredible 4-in-1 bundle. Are you interested in becoming a Python pro? Do you want to learn more about the incredible world of machine learning, and what it can do for you? Then keep reading. Created with the beginner in mind, this powerful bundle delves into the fundamentals behind Python and Machine Learning, from basic code and mathematical formulas to complex neural networks and ensemble modeling. Inside, you'll discover everything you need to know to get started with Python and Machine Learning, and begin your journey to success! In book one - MACHINE LEARNING FOR BEGINNERS, you'll learn: What is Artificial Intelligence Really, and Why is it So Powerful? Choosing the Right Kind of Machine Learning Model for You An Introduction to Statistics Reinforcement Learning and Ensemble Modeling "Random Forests" and Decision Trees In book two - MACHINE LEARNING MATHEMATICS, you will: Learn the Fundamental Concepts of Machine Learning Algorithms Understand The Four Fundamental Types of Machine Learning Algorithm Master the Concept of "Statistical Learning" Learn Everything You Need to Know about Neural Networks and Data Pipelines Master the Concept of "General Setting of Learning" In book three - LEARNING PYTHON, you'll discover: How to Install, Run, and Understand Python on Any Operating System A Comprehensive Introduction to Python Python Basics and Writing Code Writing Loops, Conditional Statements,

Exceptions and More Python Expressions and The Beauty of Inheritances And in book four - PYTHON MACHINE LEARNING, you will: Learn the Fundamentals of Machine Learning Master the Nuances of 12 of the Most Popular and Widely-Used Machine Learning Algorithms Become Familiar with Data Science Technology Dive Into the Functioning of Scikit-Learn Library and Develop Machine Learning Models Uncover the Secrets of the Most Critical Aspect of Developing a Machine Learning Model - Data Pre-Processing and Training/Testing Subsets Whether you're a complete beginner or a programmer looking to improve your skillset, this bundle is your all-in-one solution to mastering the world of Python and Machine Learning. So don't wait - it's never been easier to learn. Buy Now to Become a Master of Python and Machine Learning Today!

Deep Learning with Python

Exploring deep learning techniques and neural network architectures with PyTorch, Keras, and TensorFlow, 2nd Edition

Deep Learning with TensorFlow 2 and Keras

A Practical Guide

Probabilistic Deep Learning

AI and Machine Learning for Coders

Learning TensorFlow

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

Learn advanced state-of-the-art deep learning techniques and their applications using popular Python libraries Key Features Build a strong foundation in neural networks and deep learning with Python libraries Explore advanced deep learning techniques and their applications across computer vision and NLP Learn how a computer can navigate in complex environments with reinforcement learning Book Description With the surge in artificial intelligence in applications catering to both business and consumer needs, deep learning is more important than ever for meeting current and future market demands. With this book, you'll explore deep learning, and learn how to put machine learning to use in

your projects. This second edition of Python Deep Learning will get you up to speed with deep learning, deep neural networks, and how to train them with high-performance algorithms and popular Python frameworks. You'll uncover different neural network architectures, such as convolutional networks, recurrent neural networks, long short-term memory (LSTM) networks, and capsule networks. You'll also learn how to solve problems in the fields of computer vision, natural language processing (NLP), and speech recognition. You'll study generative model approaches such as variational autoencoders and Generative Adversarial Networks (GANs) to generate images. As you delve into newly evolved areas of reinforcement learning, you'll gain an understanding of state-of-the-art algorithms that are the main components behind popular games Go, Atari, and Dota. By the end of the book, you will be well-versed with the theory of deep learning along with its real-world applications. What you will learn

- Grasp the mathematical theory behind neural networks and deep learning processes
- Investigate and resolve computer vision challenges using convolutional networks and capsule networks
- Solve generative tasks using variational autoencoders and Generative Adversarial Networks
- Implement complex NLP tasks using recurrent networks (LSTM and GRU) and attention models
- Explore reinforcement learning and understand how agents behave in a complex environment
- Get up to date with applications of deep learning in autonomous vehicles

Who this book is for This book is for data science practitioners, machine learning engineers, and those interested in deep learning who have a basic foundation in machine learning and some Python programming experience. A background in mathematics and conceptual understanding of calculus and statistics will help you gain maximum benefit from this book.

Practical Deep Learning teaches total beginners how to build the datasets and models needed to train neural networks for your own DL projects. If you've been curious about machine learning but didn't know where to start, this is the book you've been waiting for. Focusing on the subfield of machine learning known as deep learning, it explains core concepts and gives you the foundation you need to start building your own models. Rather than simply outlining recipes for using existing toolkits, Practical Deep Learning teaches you the why of deep learning and will inspire you to explore further. All you need is basic familiarity with computer programming and high school math—the book will cover the rest. After an introduction to Python, you'll move through key topics like how to build a good training dataset, work with the scikit-learn and Keras libraries, and evaluate your models' performance. You'll also learn:

- How to use classic machine learning models like k-Nearest Neighbors, Random Forests, and Support Vector Machines
- How neural networks work and how they're trained
- How to use convolutional neural networks
- How to develop a successful deep learning model from scratch

You'll conduct experiments along the way, building to a final case study that incorporates everything you've learned. The perfect introduction to this dynamic, ever-expanding field, Practical Deep Learning will give you the skills and confidence to dive into your own machine learning projects.

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

Python Machine Learning

Ultimate Step by Step Guide to Machine Learning Using Python

Introduction to Deep Learning

TensorFlow 2.x in the Colaboratory Cloud

From Linear Regression to Reinforcement Learning

Solution for Edge Computing Applications

Build advanced computer vision applications using machine learning and deep learning techniques

Take your machine learning skills to the next level by mastering Deep Learning concepts and algorithms using Python. About This Book Explore and create intelligent systems using cutting-edge deep learning techniques Implement deep learning algorithms and work with revolutionary libraries in Python Get real-world examples and easy-to-follow tutorials on Theano, TensorFlow, H2O and more Who This Book Is For This book is for Data Science practitioners as well as aspirants who have a basic foundational understanding of Machine Learning concepts and some programming experience with Python. A mathematical background with a conceptual understanding of calculus and statistics is also desired. What You Will Learn Get a practical deep dive into deep learning algorithms Explore deep learning further with Theano, Caffe, Keras, and TensorFlow Learn about two of the most powerful techniques at the core of many practical deep learning implementations: Auto-Encoders and Restricted Boltzmann Machines Dive into Deep Belief Nets and Deep Neural Networks Discover more deep learning algorithms with Dropout and Convolutional Neural Networks Get to know device strategies so you can use deep learning algorithms and libraries in the real world In Detail With an increasing interest in AI around the world, deep learning has attracted a great deal of public attention. Every day, deep learning algorithms are used broadly across different industries. The book will give you all the practical information

available on the subject, including the best practices, using real-world use cases. You will learn to recognize and extract information to increase predictive accuracy and optimize results. Starting with a quick recap of important machine learning concepts, the book will delve straight into deep learning principles using Sci-kit learn. Moving ahead, you will learn to use the latest open source libraries such as Theano, Keras, Google's TensorFlow, and H2O. Use this guide to uncover the difficulties of pattern recognition, scaling data with greater accuracy and discussing deep learning algorithms and techniques. Whether you want to dive deeper into Deep Learning, or want to investigate how to get more out of this powerful technology, you'll find everything inside. Style and approach Python Machine Learning by example follows practical hands on approach. It walks you through the key elements of Python and its powerful machine learning libraries with the help of real world projects.

Apply neural network architectures to build state-of-the-art computer vision applications using the Python programming language
Key Features
Gain a fundamental understanding of advanced computer vision and neural network models in use today
Cover tasks such as low-level vision, image classification, and object detection
Develop deep learning models on cloud platforms and optimize them using TensorFlow Lite and the OpenVINO toolkit
Book Description Computer vision allows machines to gain human-level understanding to visualize, process, and analyze images and videos. This book focuses on using TensorFlow to help you learn advanced computer vision tasks such as image acquisition, processing, and analysis.

You'll start with the key principles of computer vision and deep learning to build a solid foundation, before covering neural network architectures and understanding how they work rather than using them as a black box. Next, you'll explore architectures such as VGG, ResNet, Inception, R-CNN, SSD, YOLO, and MobileNet. As you advance, you'll learn to use visual search methods using transfer learning. You'll also cover advanced computer vision concepts such as semantic segmentation, image inpainting with GAN's, object tracking, video segmentation, and action recognition. Later, the book focuses on how machine learning and deep learning concepts can be used to perform tasks such as edge detection and face recognition. You'll then discover how to develop powerful neural network models on your PC and on various cloud

platforms. Finally, you'll learn to perform model optimization methods to deploy models on edge devices for real-time inference. By the end of this book, you'll have a solid understanding of computer vision and be able to confidently develop models to automate tasks. What you will learn
Explore methods of feature extraction and image retrieval and visualize different layers of the neural network model
Use TensorFlow for various visual search methods for real-world scenarios
Build neural networks or adjust parameters to optimize the performance of models
Understand TensorFlow DeepLab to perform semantic segmentation on images and DCGAN for image inpainting
Evaluate your model and optimize and integrate it into your application to operate at scale
Get up to speed with techniques for performing manual and automated image annotation
Who this book is for This book is for computer vision professionals, image processing professionals, machine learning engineers and AI developers who have some knowledge of machine learning and deep learning and want to build expert-level computer vision applications. In addition to familiarity with TensorFlow, Python knowledge will be required to get started with this book.

Given the demand for AI and the ubiquity of JavaScript, TensorFlow.js was inevitable. With this Google framework, seasoned AI veterans and web developers alike can help propel the future of AI-driven websites. In this guide, author Gant Laborde--Google Developer Expert in machine learning and the web--provides a hands-on end-to-end approach to TensorFlow.js fundamentals for a broad technical audience that includes data scientists, engineers, web developers, students, and researchers. You'll begin by working through some basic examples in TensorFlow.js before diving deeper into neural network architectures, DataFrames, TensorFlow Hub, model conversion, transfer learning, and more. Once you finish this book, you'll know how to build and deploy production-ready deep learning systems with TensorFlow.js. Explore tensors, the most fundamental structure of machine learning
Convert data into tensors and back with a real-world example
Combine AI with the web using TensorFlow.js Use resources to convert, train, and manage machine learning data
Build and train your own training models from scratch
Summary Deep Learning with Python introduces the field of deep learning using the Python language and the powerful Keras library. Written by Keras creator and Google AI

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

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

Mastering Computer Vision with TensorFlow 2.x

Deep Learning with TensorFlow

Machine Learning with TensorFlow

Machine Learning

Machine Learning with TensorFlow, Second Edition

Deep Learning Classifiers with Memristive Networks

Concepts, Tools, and Techniques to Build Intelligent Systems

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).

**Start your Data Science career using Python today!*Are you ready to start your new exciting career? Ready to master artificial intelligence and deep learning concepts?Are you overwhelmed with complexity of the books on this subject?Then let this breezy and fun little book on Python, Machine Learning and Deep Learning models make you a Data Scientist in 7 days!This book continues from where the first book in the series, Ultimate Step by Step Guide to Machine Learning Using Python, left of. In the first book you were introduced to Python concepts such as: -Data Structures like Pandas -Foundational libraries like Numpy, Seaborn and Scikit-Learn-Regression analysis-Classification-Clustering-Association Learning-Dimension ReductionThis book builds on those concepts to expand on Machine Learning algorithms like: -Linear and Logistical regression-Decision tree-Support vector machines (SVM)After that, this book takes you on a journey into Deep Learning and Neural Networks with important concepts and libraries like: -Convolutional and Recurrent Neural Networks-TensorFlow-Keras-PyTorch-Keras-Apache MXNet-Microsoft Cognitive Toolkit (CNTK)The final part of the book covers all foundational concepts that are required for Amazon Web Services (AWS) Certified Machine Learning Specialization by explaining how to deploy your models at scale on Cloud technologies. While AWS is used in the book for illustrative purposes, Microsoft Azure and Google Cloud are also introduced as alternative cloud technologies. After reading this book you will be able to: -Code in Python with confidence-Build new machine learning and deep learning models from scratch-Know how to clean and prepare your data for analytics-Speak confidently about statistical analysis techniquesData Science was ranked the fast-growing field by LinkedIn and Data Scientist is one of the most highly sought after and lucrative careers in the world!If you are on the fence about making the leap to a new and lucrative career, this is the book for you!What sets this book apart from other books on the topic of Python and Machine learning: -Step by step code examples and explanation-Complex concepts*

explained visually-Real world applicability of the machine learning and deep learning models introducedWhat do I need to get started?You will have a step by step action plan in place once you finish this book and finally feel that you, can master data science and artificial intelligence and start a lucrative and rewarding career! Ready to dive in to the exciting world of Python and Deep Learning?Then scroll up to the top and hit that **BUY BUTTON!**

55% OFF for bookstores! Find Out the EASIEST Way for Mastering Machine Learning and Data Analysis in Python!

Summary Grokking Deep Learning teaches you to build deep learning neural networks from scratch! In his engaging style, seasoned deep learning expert Andrew Trask shows you the science under the hood, so you grok for yourself every detail of training neural networks. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology Deep learning, a branch of artificial intelligence, teaches computers to learn by using neural networks, technology inspired by the human brain. Online text translation, self-driving cars, personalized product recommendations, and virtual voice assistants are just a few of the exciting modern advancements possible thanks to deep learning. About the Book Grokking Deep Learning teaches you to build deep learning neural networks from scratch! In his engaging style, seasoned deep learning expert Andrew Trask shows you the science under the hood, so you grok for yourself every detail of training neural networks. Using only Python and its math-supporting library, NumPy, you'll train your own neural networks to see and understand images, translate text into different languages, and even write like Shakespeare! When you're done, you'll be fully prepared to move on to mastering deep learning frameworks. What's inside The science behind deep learning Building and training your own neural networks Privacy concepts, including federated learning Tips for continuing your pursuit of deep learning About the Reader For readers with high school-level math and intermediate programming skills. About the Author Andrew Trask is a PhD student at Oxford University and a research scientist at DeepMind. Previously, Andrew was a researcher and analytics product manager at Digital Reasoning, where he trained the world's largest artificial neural network and helped guide the analytics roadmap for the Synthesys cognitive computing platform. Table of Contents Introducing deep learning: why you should learn it Fundamental concepts: how do machines learn? Introduction to neural prediction: forward propagation Introduction to neural learning: gradient descent Learning multiple weights at a time: generalizing gradient descent Building your first deep neural network: introduction to backpropagation How to picture neural networks: in your head and on paper Learning signal and ignoring noise:introduction to regularization and batching Modeling probabilities and nonlinearities: activation functions Neural learning about edges and corners: intro to convolutional neural networks Neural networks that understand language: king - man + woman == ? Neural networks that write like Shakespeare: recurrent layers for variable-length data Introducing automatic optimization: let's build a deep learning framework Learning to write like Shakespeare: long short-term memory Deep learning on unseen data: introducing federated learning Where to go from here: a brief guide

Programming with TensorFlow

Deep Learning Models for Research and Industry

A Beginner's Guide to Python Programming for Machine Learning and Deep Learning, Data Analysis, Algorithms and Data Science With Scikit Learn, TensorFlow, PyTorch and Keras.

A Python-Based Introduction

Introduction to Deep Learning (Black/White Version)
Deep Learning Crash Course for Beginners with Python
Tensorflow Machine Learning

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.

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

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

Introduction to Deep Learning and Neural Networks with Python™: A Practical Guide is an intensive step-by-step guide for neuroscientists to fully understand, practice, and build neural networks. Providing math and Python™ code examples to clarify neural network calculations, by book's end readers will fully understand how neural networks work starting from the simplest model $Y=X$ and building from scratch. Details and explanations are provided on how a generic gradient descent algorithm works based on mathematical and Python™ examples, teaching you how to use the gradient descent algorithm to manually perform all calculations in both the forward and backward passes of training a neural network. Examines the practical side of deep learning and neural networks Provides a problem-based approach to building artificial neural networks using real data Describes Python™ functions and features for neuroscientists Uses a careful tutorial approach to describe implementation of neural networks in Python™ Features math and code examples (via companion website) with helpful instructions for easy implementation

Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow

Ultimate Step by Step Guide to Deep Learning Using Python

Develop Deep Learning Models on Theano and TensorFlow Using Keras

An Introduction to Deep Learning on Google's Cloud Service

Theory and Applications of Artificial Neural Networks, CNN, RNN, LSTM and Autoencoders Using TensorFlow 2.0- Contains

Exercises with Solutions and Hands-On Projects

A Comprehensive Beginner's Guide to Learn About the Realms of TensorFlow Machine Learning From A-Z

Python Deep Learning

****Start your Data Science career using Python today!* Are you ready to start your new exciting career? Ready to crush your machine learning career goals? Are you overwhelmed with complexity of the books on this subject? Then let this breezy and fun little book on Python and machine learning models make you a data scientist in 7 days! First part of this book introduces Python basics including: 1) Data Structures like Pandas 2) Foundational libraries like Numpy, Seaborn and Scikit-Learn Second part of this book shows you how to build predictive machine learning models step by step using techniques such as: 1) Regression analysis 2) Decision tree analysis 3) Training and testing data models 4) And much more! After reading this book you will be able to: 1) Code in Python with confidence 2) Build new machine learning models from scratch 3) Know how to clean and prepare your data for analytics 4) Speak confidently about statistical analysis techniques Data Science was ranked the fast-growing field by LinkedIn and Data Scientist is one of the most highly sought after and lucrative careers in the world! If you are on the fence about making the leap to a new and lucrative career, this is the book for you! What sets this book apart from other books on the topic of Python and Machine learning: 1) Step by step code examples and explanation 2) Complex concepts explained visually 3) Real world applicability of the machine learning models introduced 4) Bonus free code samples that you can try yourself without any prior experience in Python! What do I need to get started? You will have a step by step action plan in place once you finish this book and finally feel that you, can master data science and machine learning and start lucrative and rewarding career! Ready to dive in to the exciting world of Python and Machine Learning? Then scroll up to the top and hit that BUY BUTTON!***

If you're looking to make a career move from programmer to AI specialist, this is the ideal place to start. Based on Laurence Moroney's extremely successful AI courses, this introductory book provides a hands-on, code-first approach to help you build confidence while you learn key topics. You'll understand how to implement the most common scenarios in machine learning, such as computer vision, natural language processing (NLP), and sequence modeling for web, mobile, cloud, and embedded runtimes. Most books on machine learning begin with a daunting amount of advanced math. This guide is built on practical lessons that let you work directly with the code. You'll learn: How to build models with TensorFlow using skills that employers desire The basics of machine learning by working with code samples How to implement computer vision, including feature detection in images How to use NLP to tokenize and sequence words and sentences Methods for embedding models in Android and iOS How to serve models over the web and in the cloud with TensorFlow

Serving

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

Unlock deeper insights into Machine Learning with this vital guide to cutting-edge predictive analytics About This Book Leverage Python's most powerful open-source libraries for deep learning, data wrangling, and data visualization Learn effective strategies and best practices to improve and optimize machine learning systems and algorithms Ask - and answer - tough questions of your data with robust statistical models, built for a range of datasets Who This Book Is For If you want to find out how to use Python to start answering critical questions of your data, pick up Python Machine Learning - whether you want to get started from scratch or want to extend your data science knowledge, this is an essential and unmissable resource. What You Will Learn Explore how to use different machine learning models to ask different questions of your data Learn how to build neural networks using Keras and Theano Find out how to write clean and elegant Python code that will optimize the strength of your algorithms Discover how to embed your machine learning model in a web application for increased accessibility Predict continuous target outcomes using regression analysis Uncover hidden patterns and structures in data with clustering Organize data using effective pre-processing techniques Get to grips with sentiment analysis to delve deeper into textual and social media data In Detail Machine learning and predictive analytics are transforming the way businesses and other organizations operate. Being able to understand trends and patterns in complex data is critical to success, becoming one of the key strategies for unlocking growth in a challenging contemporary marketplace. Python can help you deliver key insights into your data - its unique capabilities as a language let you build sophisticated

algorithms and statistical models that can reveal new perspectives and answer key questions that are vital for success. Python Machine Learning gives you access to the world of predictive analytics and demonstrates why Python is one of the world's leading data science languages. If you want to ask better questions of data, or need to improve and extend the capabilities of your machine learning systems, this practical data science book is invaluable. Covering a wide range of powerful Python libraries, including scikit-learn, Theano, and Keras, and featuring guidance and tips on everything from sentiment analysis to neural networks, you'll soon be able to answer some of the most important questions facing you and your organization. Style and approach Python Machine Learning connects the fundamental theoretical principles behind machine learning to their practical application in a way that focuses you on asking and answering the right questions. It walks you through the key elements of Python and its powerful machine learning libraries, while demonstrating how to get to grips with a range of statistical models.

Artificial Intelligence and Neural Network Concepts Explained in Simple Terms

Practical Deep Learning

Implement Machine Learning and Deep Learning Models with Python

An Introduction to Deep Learning on Google's Cloud Service

Get up to speed with the newly introduced features of TensorFlow 2.0

A Guide to Building Deep Learning Systems

Deep Learning for Coders with fastai and PyTorch

Are you interested in learning machine learning and deep learning? TensorFlow is the single most popular library available today. Offering some of the very best graph computations, TensorFlow helps data scientists in designing neural networks using a cool feature called TensorBoard. It has support for both recurrent neural networks (RNNs) and convolution, as well as parallel processing support on GPU and CPU. While TensorFlow is an incredibly important machine and deep learning library, we also give you an introduction to three others - NumPy, Pandas, and Scikit Learn. I have produced a hands-on guide, with plenty of code examples for you to follow along with. Here's what you will learn: -What deep learning is-The difference between deep learning and machine learning-What TensorFlow is-How to install it on Windows and Mac-The basics of TensorFlow-Using TensorBoard-About NumPy, Scikit Learn, and Pandas-About linear regression-Kernel methods-Building an Artificial Neural Network using TensorFlow-TensorFlow image classification-TensorFlow autoencoders-Much more! If you are already proficient at programming in Python and are ready to take the next step into machine learning, this guide is for you. Scroll up, hit that Buy Now button, and set off on a brand new machine learning journey.

About the book: In Computer Sciences there is currently a gold rush mood due to a new field called "Deep Learning". But what is Deep Learning? This book is an introduction to Neural Networks and the most important Deep Learning model - the Convolutional Neural Network model including a description of tricks that can be used to train such models more quickly. We start with the biological role model: the Neuron. About 86.000.000.000 of these simple processing elements are in your brain! And they all work in parallel! We discuss how to model the operation of a biological neuron with technical neuron models and then consider the first simple single-layer network of technical neurons. We then introduce the Multi-Layer Perceptron (MLP) and the Convolutional Neural Network (CNN) model which uses the MLP at its end. At the end of the book we discuss promising new directions for the field of Deep Learning. A famous physicist once said: "What I cannot create, I do not understand". For this, the book is full of examples of how to program all models discussed in Python and TensorFlow - Today, the most important Deep Learning library. About the author: Prof. Dr.-Ing. Juergen Brauer is a professor for Sensor Data Processing and Programming at the University of Applied Sciences Kempten in Germany where he holds a "Deep Learning" and other machine learning related lectures for Computer Science and Advanced Driver Assistance Systems students. His personal experience tells him: "What I cannot program, I do not understand". Note: This is a black/white version of the original book version which uses colors in some figures. It allows for lower printing costs and therefore a lower final price. However, some details might be recognized better in the color version.

Updated with new code, new projects, and new chapters, Machine Learning with TensorFlow, Second Edition gives readers a solid foundation in machine-learning concepts and the TensorFlow library. Summary Updated with new code, new projects, and new chapters, Machine Learning with TensorFlow, Second Edition gives readers a solid foundation in machine-learning concepts and the TensorFlow library. Written by NASA JPL Deputy CTO and Principal Data Scientist Chris Mattmann, all examples are accompanied by downloadable Jupyter Notebooks for a hands-on experience coding TensorFlow with Python. New and revised content expands coverage of core machine learning algorithms, and advancements in neural networks such as VGG-Face facial identification classifiers and deep speech classifiers. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the technology Supercharge your data analysis with machine learning! ML algorithms automatically improve as they process data, so results get better over

time. You don't have to be a mathematician to use ML: Tools like Google's TensorFlow library help with complex calculations so you can focus on getting the answers you need. About the book *Machine Learning with TensorFlow, Second Edition* is a fully revised guide to building machine learning models using Python and TensorFlow. You'll apply core ML concepts to real-world challenges, such as sentiment analysis, text classification, and image recognition. Hands-on examples illustrate neural network techniques for deep speech processing, facial identification, and auto-encoding with CIFAR-10. What's inside *Machine Learning with TensorFlow* Choosing the best ML approaches Visualizing algorithms with TensorBoard Sharing results with collaborators Running models in Docker About the reader Requires intermediate Python skills and knowledge of general algebraic concepts like vectors and matrices. Examples use the super-stable 1.15.x branch of TensorFlow and TensorFlow 2.x. About the author Chris Mattmann is the Division Manager of the Artificial Intelligence, Analytics, and Innovation Organization at NASA Jet Propulsion Lab. The first edition of this book was written by Nishant Shukla with Kenneth Fricklas. Table of Contents PART 1 - YOUR MACHINE-LEARNING RIG 1 A machine-learning odyssey 2 TensorFlow essentials PART 2 - CORE LEARNING ALGORITHMS 3 Linear regression and beyond 4 Using regression for call-center volume prediction 5 A gentle introduction to classification 6 Sentiment classification: Large movie-review dataset 7 Automatically clustering data 8 Inferring user activity from Android accelerometer data 9 Hidden Markov models 10 Part-of-speech tagging and word-sense disambiguation PART 3 - THE NEURAL NETWORK PARADIGM 11 A peek into autoencoders 12 Applying autoencoders: The CIFAR-10 image dataset 13 Reinforcement learning 14 Convolutional neural networks 15 Building a real-world CNN: VGG-Face and 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 introduces readers to the fundamentals of deep neural network architectures, with a special emphasis on memristor circuits and systems. At first, the book offers an overview of neuro-memristive systems, including memristor devices, models, and theory, as well as an introduction to deep learning neural networks such as multi-layer networks, convolution neural networks, hierarchical temporal memory, and long short term memories, and deep neuro-fuzzy networks. It then focuses on the design of these neural networks using memristor crossbar architectures in detail. The book integrates the theory with various applications of neuro-memristive circuits and systems. It provides an introductory tutorial on a range of issues in the design, evaluation techniques, and implementations of different deep neural

network architectures with memristors.

Artificial Intelligence with Python

Probabilistic Programming and Bayesian Inference

4 Books in 1: A Complete Overview for Beginners to Master the Basics of Python Programming and Understand How to Build Artificial Intelligence Through Data Science

A Hands-on Introduction

Theory and Applications

Neural networks in TensorFlow.js

Deep Learning With Python

Deep learning is often viewed as the exclusive domain of math PhDs and big tech companies. But as this hands-on guide demonstrates, programmers comfortable with Python can achieve impressive results in deep learning with little math background, small amounts of data, and minimal code. How? With fastai, the first library to provide a consistent interface to the most frequently used deep learning applications. Authors Jeremy Howard and Sylvain Gugger, the creators of fastai, show you how to train a model on a wide range of tasks using fastai and PyTorch. You ' ll also dive progressively further into deep learning theory to gain a complete understanding of the algorithms behind the scenes. Train models in computer vision, natural language processing, tabular data, and collaborative filtering Learn the latest deep learning techniques that matter most in practice Improve accuracy, speed, and reliability by understanding how deep learning models work Discover how to turn your models into web applications Implement deep learning algorithms from scratch Consider the ethical implications of your work Gain insight from the foreword by PyTorch cofounder, Soumith Chintala

Artificial intelligence is the rage today! While you may find it difficult to understand the most recent advancements in AI, it simply boils down to two most celebrated developments: Machine Learning and Deep Learning. In 2020, Deep Learning is leagues ahead because of its supremacy when it comes to accuracy, especially when trained with enormous amounts of data. Deep Learning, essentially, is a subset of Machine Learning, but it's capable of achieving tremendous power and flexibility. And the era of big data technology presents vast opportunities for incredible innovations in deep learning. How Is This Book Different? This book gives equal importance to the theoretical as well as practical aspects of deep learning. You will understand how high-performing deep learning algorithms work. In every chapter, the theoretical explanation of the different types of deep learning techniques is followed by practical examples. You will learn how to implement different deep learning techniques using the TensorFlow Keras library for Python. Each chapter contains exercises that you can use to assess your understanding of the concepts explained in that chapter. Also, in the Resources, the Python notebook for each chapter is provided. The key advantage of buying this book is you get instant access to all the extra content presented with this book--Python codes, references, exercises, and PDFs--on the publisher's website. You don't need to spend an extra cent. The datasets used in this book are either downloaded at runtime or are available in the Resources/Datasets folder. Another advantage is a detailed explanation of the installation steps for the software that you will need to implement the various deep learning algorithms in this book is provided. That is, you get to experiment with the practical aspects of Deep Learning right from page 1. Even if you are new to Python, you will find the crash

course on Python programming language in the first chapter immensely useful. Since all the codes and datasets are included with this book, you only need access to a computer with the internet to get started. The topics covered include: Python Crash Course Deep Learning Prerequisites: Linear and Logistic Regression Neural Networks from Scratch in Python Introduction to TensorFlow and Keras Convolutional Neural Networks Sequence Classification with Recurrent Neural Networks Deep Learning for Natural Language Processing Unsupervised Learning with Autoencoders Answers to All Exercises Click the BUY button and download the book now to start your Deep Learning journey.

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 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 D ü rr 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 Master Bayesian Inference through Practical Examples and Computation – Without Advanced Mathematical Analysis Bayesian methods of inference are deeply natural and extremely powerful. However, most discussions of Bayesian inference rely on intensely complex mathematical analyses and artificial examples, making it inaccessible to anyone without a strong mathematical background. Now, though,

Cameron Davidson-Pilon introduces Bayesian inference from a computational perspective, bridging theory to practice – freeing you to get results using computing power. Bayesian Methods for Hackers illuminates Bayesian inference through probabilistic programming with the powerful PyMC language and the closely related Python tools NumPy, SciPy, and Matplotlib. Using this approach, you can reach effective solutions in small increments, without extensive mathematical intervention. Davidson-Pilon begins by introducing the concepts underlying Bayesian inference, comparing it with other techniques and guiding you through building and training your first Bayesian model. Next, he introduces PyMC through a series of detailed examples and intuitive explanations that have been refined after extensive user feedback. You ' ll learn how to use the Markov Chain Monte Carlo algorithm, choose appropriate sample sizes and priors, work with loss functions, and apply Bayesian inference in domains ranging from finance to marketing. Once you ' ve mastered these techniques, you ' ll constantly turn to this guide for the working PyMC code you need to jumpstart future projects. Coverage includes

- Learning the Bayesian “ state of mind ” and its practical implications
- Understanding how computers perform Bayesian inference
- Using the PyMC Python library to program Bayesian analyses
- Building and debugging models with PyMC
- Testing your model ' s “ goodness of fit ”
- Opening the “ black box ” of the Markov Chain Monte Carlo algorithm to see how and why it works
- Leveraging the power of the “ Law of Large Numbers ”
- Mastering key concepts, such as clustering, convergence, autocorrelation, and thinning
- Using loss functions to measure an estimate ' s weaknesses based on your goals and desired outcomes
- Selecting appropriate priors and understanding how their influence changes with dataset size
- Overcoming the “ exploration versus exploitation ” dilemma: deciding when “ pretty good ” is good enough
- Using Bayesian inference to improve A/B testing
- Solving data science problems when only small amounts of data are available

Cameron Davidson-Pilon has worked in many areas of applied mathematics, from the evolutionary dynamics of genes and diseases to stochastic modeling of financial prices. His contributions to the open source community include lifelines, an implementation of survival analysis in Python. Educated at the University of Waterloo and at the Independent University of Moscow, he currently works with the online commerce leader Shopify.

TensorFlow for Deep Learning

Introduction to Deep Learning and Neural Networks with PythonTM

TensorFlow 2.0 Quick Start Guide

Deep Learning with JavaScript

With Python, Keras and TensorFlow Probability

Machine Learning for Economics and Finance in TensorFlow 2

With Complete Python and TensorFlow 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.

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.

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

Use TensorFlow 2.x with Google's Colaboratory (Colab) product that

offers a free cloud service for Python programmers. Colab is especially well suited as a platform for TensorFlow 2.x deep learning applications. You will learn Colab's default install of the most current TensorFlow 2.x along with Colab's easy access to on-demand GPU hardware acceleration in the cloud for fast execution of deep learning models. This book offers you the opportunity to grasp deep learning in an applied manner with the only requirement being an Internet connection. Everything else—Python, TensorFlow 2.x, GPU support, and Jupyter Notebooks—is provided and ready to go from Colab. The book begins with an introduction to TensorFlow 2.x and the Google Colab cloud service. You will learn how to provision a workspace on Google Colab and build a simple neural network application. From there you will progress into TensorFlow datasets and building input pipelines in support of modeling and testing. You will find coverage of deep learning classification and regression, with clear code examples showing how to perform each of those functions. Advanced topics covered in the book include convolutional neural networks and recurrent neural networks. This book contains all the applied math and programming you need to master the content. Examples range from simple to relatively complex when necessary to ensure acquisition of appropriate deep learning concepts and constructs. Examples are carefully explained, concise, accurate, and complete to perfectly complement deep learning skill development. Care is taken to walk you through the foundational principles of deep learning through clear examples written in Python that you can try out and experiment with using Google Colab from the comfort of your own home or office. What You Will Learn Be familiar with the basic concepts and constructs of applied deep learning Create machine learning models with clean and reliable Python code Work with datasets common to deep learning applications Prepare data for TensorFlow consumption Take advantage of Google Colab's built-in support for deep learning Execute deep learning experiments using a variety of neural network models Be able to mount Google Colab directly to your Google Drive account Visualize training versus test performance to see model fit Who This Book Is For Readers who want to learn the highly popular TensorFlow 2.x deep learning platform, those who wish to master deep learning fundamentals that are sometimes skipped over in the rush to be productive, and those looking to build competency with a modern cloud service tool such as Google Colab

Regression, ConvNets, GANs, RNNs, NLP, and more with TensorFlow 2 and the Keras API, 2nd Edition

Grokking Deep Learning

Bayesian Methods for Hackers

Predictive Modelling Concepts Explained in Simple Terms for Beginners

Learn TensorFlow 2.0

Machine Learning and Deep Learning Using Python and TensorFlow

Learning TensorFlow.js

Discover the practical aspects of implementing deep-learning solutions using the rich Python ecosystem. This book bridges the gap between the academic state-of-

the-art and the industry state-of-the-practice by introducing you to deep learning frameworks such as Keras, Theano, and Caffe. The practicalities of these frameworks is often acquired by practitioners by reading source code, manuals, and posting questions on community forums, which tends to be a slow and a painful process. Deep Learning with Python allows you to ramp up to such practical know-how in a short period of time and focus more on the domain, models, and algorithms. This book briefly covers the mathematical prerequisites and fundamentals of deep learning, making this book a good starting point for software developers who want to get started in deep learning. A brief survey of deep learning architectures is also included. Deep Learning with Python also introduces you to key concepts of automatic differentiation and GPU computation which, while not central to deep learning, are critical when it comes to conducting large scale experiments. What You Will Learn Leverage deep learning frameworks in Python namely, Keras, Theano, and Caffe Gain the fundamentals of deep learning with mathematical prerequisites Discover the practical considerations of large scale experiments Take deep learning models to production Who This Book Is For Software developers who want to try out deep learning as a practical solution to a particular problem. Software developers in a data science team who want to take deep learning models developed by data scientists to production.

Use TensorFlow 2.x with Google's Colaboratory (Colab) product that offers a free cloud service for Python programmers. Colab is especially well suited as a platform for TensorFlow 2.x deep learning applications. You will learn Colab's default install of the most current TensorFlow 2.x along with Colab's easy access to on-demand GPU hardware acceleration in the cloud for fast execution of deep learning models. This book offers you the opportunity to grasp deep learning in an applied manner with the only requirement being an Internet connection. Everything else--Python, TensorFlow 2.x, GPU support, and Jupyter Notebooks--is provided and ready to go from Colab. The book begins with an introduction to TensorFlow 2.x and the Google Colab cloud service. You will learn how to provision a workspace on Google Colab and build a simple neural network application. From there you will progress into TensorFlow datasets and building input pipelines in support of modeling and testing. You will find coverage of deep learning classification and regression, with clear code examples showing how to perform each of those functions. Advanced topics covered in the book include convolutional neural networks and recurrent neural networks. This book contains all the applied math and programming you need to master the content. Examples range from simple to relatively complex when necessary to ensure acquisition of appropriate deep learning concepts and constructs. Examples are carefully explained, concise, accurate, and complete to perfectly complement deep learning skill development. Care is taken to walk you through the foundational principles of deep learning through clear examples written in Python that you can try out and experiment with using Google Colab from the comfort of your own

home or office. You will: Be familiar with the basic concepts and constructs of applied deep learning Create machine learning models with clean and reliable Python code Work with datasets common to deep learning applications Prepare data for TensorFlow consumption Take advantage of Google Colab's built-in support for deep learning Execute deep learning experiments using a variety of neural network models Be able to mount Google Colab directly to your Google Drive account Visualize training versus test performance to see model fit. Work on economic problems and solutions with tools from machine learning. ML has taken time to move into the space of academic economics. This is because empirical work in economics is concentrated on the identification of causal relationships in parsimonious statistical models; whereas machine learning is oriented towards prediction and is generally uninterested in either causality or parsimony. That leaves a gap for both students and professionals in the economics industry without a standard reference. This book focuses on economic problems with an empirical dimension, where machine learning methods may offer something of value. This includes coverage of a variety of discriminative deep learning models (DNNs, CNNs, RNNs, LSTMs, the Transformer Model, etc.), generative machine learning models, random forests, gradient boosting, clustering, and feature extraction. You'll also learn about the intersection of empirical methods in economics and machine learning, including regression analysis, text analysis, and dimensionality reduction methods, such as principal components analysis. TensorFlow offers a toolset that can be used to setup and solve any mathematical model, including those commonly used in economics. This book is structured to teach through a sequence of complete examples, each framed in terms of a specific economic problem of interest or topic. Otherwise complicated content is then distilled into accessible examples, so you can use TensorFlow to solve workhorse models in economics and finance. What You'll Learn Define, train, and evaluate machine learning models in TensorFlow 2 Apply fundamental concepts in machine learning, such as deep learning and natural language processing, to economic and financial problems Solve workhorse models in economics and finance Who This Book Is For Students and data scientists working in the economics industry. Academic economists and social scientists who have an interest in machine learning are also likely to find this book useful.

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.