

Neural Network Programming With Python Create Your Own Neural Network

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.

This book doesn't have any superpowers or magic formula to help you master the art of neural networks and deep learning. We believe that such learning is all in your heart. You need to learn a concept by heart and then brainstorm its different possibilities. I don't claim that after reading this book you will become an expert in Python and Deep Learning Neural Networks. Instead, you will, for sure, have a basic understanding of deep learning and its implications and real-life applications. Most of the time, what confuses us is the application of a certain thing in our lives. Once we know that, we can relate the subject to that particular thing and learn. An interesting thing is that neural networks also learn the same way. This makes it easier to learn about them when we know the basics. Let's take a look at what this book has to offer: ? The basics of Python including data types, operators and numbers. ? Advanced programming in Python with Python expressions, types and much more. ? A comprehensive overview of deep learning and its link to the smart systems that we are now building. ? An overview of how artificial neural networks work in real life. ? An overview of PyTorch. ? An overview of TensorFlow. ? An overview of Keras. ? How to create a convolutional neural network. ? A comprehensive understanding of deep learning applications and its ethical implications, including in the present and future. This book offers you the basic knowledge about Python and Deep Learning Neural Networks that you will need to lay the foundation for future studies. This book will start you on the road to mastering the art of deep learning neural networks. When I say that I don't have the magic formula to make you learn, I mean it. My point is that you should learn Python coding and Python libraries to build neural networks by practicing hard. The more you practice, the better it is for your skills. It is only after thorough and in depth practice that you will be able to create your own programs. Unlike other books, I don't claim that this book will make you a master of deep learning after a single read. That's not realistic, in fact, it's even a bit absurd. What I claim is that you will definitely learn about the basics. The rest is practice. The more you practice the better you code.

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.

This second edition of Foundations of Python Network Programming targets Python 2.5 through Python 2.7, the most popular production versions of the language. Python has made great strides since Apress released the first edition of this book back in the days of Python 2.3. The advances required new chapters to be written from the ground up, and others to be extensively revised. You will learn fundamentals like IP, TCP, DNS and SSL by using working Python programs; you will also be able to familiarize yourself with infrastructure components like memcached and message queues. You can also delve into network server designs, and compare threaded approaches with asynchronous event-based solutions. But the biggest change is this edition's expanded treatment of the web. The HTTP protocol is covered in extensive detail, with each feature accompanied by sample Python code. You can use your HTTP protocol expertise by studying an entire chapter on screen scraping and you can then test lxml and BeautifulSoup against a real-world web site. The chapter on web application programming now covers both the WSGI standard for component interoperability, as well as modern web frameworks like Django. Finally, all of the old favorites from the first edition are back: E-mail protocols like SMTP, POP, and IMAP get full treatment, as does XML-RPC. You can still learn how to code Python network programs using the Telnet and FTP protocols, but you are likely to appreciate the power of more modern alternatives like the paramiko SSH2 library. If you are a Python programmer who needs to learn the network, this is the book that you want by your side.

??Buy the Paperback Version of this Book and get the Kindle Book version for FREE ?? Step into the fascinating world of data science.. You to participate in the revolution that brings artificial intelligence back to the heart of our society, thanks to data scientists. Data science consists in translating problems of any other nature into quantitative modeling problems, solved by processing algorithms. This book, designed for anyone wishing to learn Deep Learning. This book presents the main techniques: deep neural networks, able to model all kinds of data, convolution networks, able to classify images, segment them and discover the objects or people who are there, recurring networks, it contains sample code so that the reader can easily test and run the programs. On the program: Deep learning Neural Networks and Deep Learning Deep Learning Parameters and Hyper-parameters Deep Neural Networks Layers Deep Learning Activation Functions Convolutional Neural Network Python Data Structures Best practices in Python and Zen of Python Installing Python Python These are some of the topics covered in this book: fundamentals of deep learning fundamentals of probability fundamentals of statistics fundamentals of linear algebra introduction to machine learning and deep learning fundamentals of machine learning fundamentals of neural networks and deep learning deep learning parameters and hyper-parameters deep neural networks layers deep learning activation functions convolutional neural network Deep learning in practice (in jupyter notebooks) python data structures best practices in python and zen of python installing python The following are the objectives of this book: To help you understand deep learning in detail To help you know how to get started with deep learning in Python by setting up the coding environment. To help you transition from a deep learning Beginner to a Professional. To help you learn how to develop a complete and functional artificial neural network model in Python on your own. And more Get this book now to learn more about -- Deep learning in Python by setting up the coding environment.!

Learn how to develop intelligent applications with sequential learning and apply modern methods for language modeling with neural network architectures for deep learning with Python's most popular

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

Build smarter programs with the power of neural networks and the simplicity of Python About This Book* Make your roots stronger in neural networks by this concept-rich yet highly practical guide; from single layer to multiple layers with the help of Python* Through this book, you will develop a strong background in neural networks, regardless of your level of previous knowledge in this subject* You will be able to implement solutions from scratch, so the whole process on foundations of neural network solution design will be paced by you Who This Book Is For This book is designed for novices as well as intermediate Python developers who have a statistical background and want to work with neural networks to get better results from complex data. It also contains enough food for thought for those who want to improve their skills in machine learning and deep learning. What You Will Learn* See the latest innovations in the field* Become fluent in Python to develop neural networks solutions capable of solving complex and interesting tasks* Implement neural networks step-by-step* Solve your complex computational problems with the aid of neural networks and Python* The reader will be able to set up his/her neural network with ease, according to the objective he/she wants to apply.* The reader will be able to design time series based models using RNNs in Python.* Will be able to design high level solutions with CNNs in Python In Detail If you wish to solve your complex computational problem efficiently, neural networks come to the rescue. This book will teach you how to ace neural networks and solve your computational problems with Python—right from predicting to self-learning models—with ease. We start off with neural network design, then you'll build a solid foundational knowledge of how a neural network learns from data, and the principles behind it. This book covers various types of neural networks including recurrent neural networks and convoluted neural networks. You will not only learn how to train neural networks, but also see a generalization of these networks. With the help of practical examples and real-world use cases, you will learn to implement these neural networks in your applications.

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.

"We finally have the definitive treatise on PyTorch! It covers the basics and abstractions in great detail. I hope this book becomes your extended reference document." —Soumith Chintala, co-creator of PyTorch Key Features Written by PyTorch's creator and key contributors Develop deep learning models in a familiar Pythonic way Use PyTorch to build an image classifier for cancer detection Diagnose problems with your neural network and improve training with data augmentation Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About The Book 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. Instantly familiar to anyone who knows Python data tools like NumPy and Scikit-learn, PyTorch simplifies deep learning without sacrificing advanced features. It's great for building quick models, and it scales smoothly from laptop to enterprise. Deep Learning with PyTorch teaches you to create deep learning and neural network systems with PyTorch. This practical book gets you to work right away building a tumor image classifier from scratch. After covering the basics, you'll learn best practices for the entire deep learning pipeline, tackling advanced projects as your PyTorch skills become more sophisticated. All code samples are easy to explore in downloadable Jupyter notebooks. What You Will Learn Understanding deep learning data structures such as tensors and neural networks Best practices for the PyTorch Tensor API, loading data in Python, and visualizing results Implementing modules and loss functions Utilizing pretrained models from PyTorch Hub Methods for training networks with limited inputs Sifting through unreliable results to diagnose and fix problems in your neural network Improve your results with augmented data, better model architecture, and fine tuning This Book Is Written For For Python programmers with an interest in machine learning. No experience with PyTorch or other deep learning frameworks is required. About The Authors Eli Stevens has worked in Silicon Valley for the past 15 years as a software engineer, and the past 7 years as Chief Technical Officer of a startup making medical device software. Luca Antiga is co-founder and CEO of an AI engineering company located in Bergamo, Italy, and a regular contributor to PyTorch. Thomas Viehmann is a Machine Learning and PyTorch speciality trainer and 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

You've decided to tackle machine learning - because you're job hunting, embarking on a new project, or just think self-driving cars are cool. But where to start? It's easy to be intimidated, even as a software developer. The good news is that it doesn't have to be that hard. Master machine learning by writing code one line at a time, from simple learning programs all the way to a true deep learning system. Tackle

the hard topics by breaking them down so they're easier to understand, and build your confidence by getting your hands dirty. Peel away the obscurities of machine learning, starting from scratch and going all the way to deep learning. Machine learning can be intimidating, with its reliance on math and algorithms that most programmers don't encounter in their regular work. Take a hands-on approach, writing the Python code yourself, without any libraries to obscure what's really going on. Iterate on your design, and add layers of complexity as you go. Build an image recognition application from scratch with supervised learning. Predict the future with linear regression. Dive into gradient descent, a fundamental algorithm that drives most of machine learning. Create perceptrons to classify data. Build neural networks to tackle more complex and sophisticated data sets. Train and refine those networks with backpropagation and batching. Layer the neural networks, eliminate overfitting, and add convolution to transform your neural network into a true deep learning system. Start from the beginning and code your way to machine learning mastery. What You Need: The examples in this book are written in Python, but don't worry if you don't know this language: you'll pick up all the Python you need very quickly. Apart from that, you'll only need your computer, and your code-adept brain.

Learn, understand, and implement deep neural networks in a math- and programming-friendly approach using Keras and Python. The book focuses on an end-to-end approach to developing supervised learning algorithms in regression and classification with practical business-centric use-cases implemented in Keras. The overall book comprises three sections with two chapters in each section. The first section prepares you with all the necessary basics to get started in deep learning. Chapter 1 introduces you to the world of deep learning and its difference from machine learning, the choices of frameworks for deep learning, and the Keras ecosystem. You will cover a real-life business problem that can be solved by supervised learning algorithms with deep neural networks. You'll tackle one use case for regression and another for classification leveraging popular Kaggle datasets. Later, you will see an interesting and challenging part of deep learning: hyperparameter tuning; helping you further improve your models when building robust deep learning applications. Finally, you'll further hone your skills in deep learning and cover areas of active development and research in deep learning. At the end of Learn Keras for Deep Neural Networks, you will have a thorough understanding of deep learning principles and have practical hands-on experience in developing enterprise-grade deep learning solutions in Keras. What You'll Learn Master fast-paced practical deep learning concepts with math- and programming-friendly abstractions. Design, develop, train, validate, and deploy deep neural networks using the Keras framework Use best practices for debugging and validating deep learning models Deploy and integrate deep learning as a service into a larger software service or product Extend deep learning principles into other popular frameworks Who This Book Is For Software engineers and data engineers with basic programming skills in any language and who are keen on exploring deep learning for a career move or an enterprise project.

Design and create neural networks with deep learning and artificial intelligence principles using OpenAI Gym, TensorFlow, and Keras Key Features Explore neural network architecture and understand how it functions Learn algorithms to solve common problems using back propagation and perceptrons Understand how to apply neural networks to applications with the help of useful illustrations Book Description Neural networks play a very important role in deep learning and artificial intelligence (AI), with applications in a wide variety of domains, right from medical diagnosis, to financial forecasting, and even machine diagnostics. Hands-On Neural Networks is designed to guide you through learning about neural networks in a practical way. The book will get you started by giving you a brief introduction to perceptron networks. You will then gain insights into machine learning and also understand what the future of AI could look like. Next, you will study how embeddings can be used to process textual data and the role of long short-term memory networks (LSTMs) in helping you solve common natural language processing (NLP) problems. The later chapters will demonstrate how you can implement advanced concepts including transfer learning, generative adversarial networks (GANs), autoencoders, and reinforcement learning. Finally, you can look forward to further content on the latest advancements in the field of neural networks. By the end of this book, you will have the skills you need to build, train, and optimize your own neural network model that can be used to provide predictable solutions. What you will learn Learn how to train a network by using backpropagation Discover how to load and transform images for use in neural networks Study how neural networks can be applied to a varied set of applications Solve common challenges faced in neural network development Understand the transfer learning concept to solve tasks using Keras and Visual Geometry Group (VGG) network Get up to speed with advanced and complex deep learning concepts like LSTMs and NLP Explore innovative algorithms like GANs and deep reinforcement learning Who this book is for If you are interested in artificial intelligence and deep learning and want to further your skills, then this intermediate-level book is for you. Some knowledge of statistics will help you get the most out of this book.

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.

Neural Network Programming With Python Create Your Own Neural Network!

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

Programming With Python - 8 BOOK BUNDLE!! Deep Learning With Keras Here Is A Preview Of What You'll Learn Here... The difference between deep learning and machine learning Deep neural networks Convolutional neural networks Building deep learning models with Keras Multi-layer perceptron network models And much more... Convolutional Neural Networks In Python Here Is A Preview Of What You'll Learn Here... Convolutional neural networks structure How convolutional neural networks actually work Convolutional neural networks applications The importance of convolution operator How to build a simple image classification CNN And much, much more! Python Machine Learning Here Is A Preview Of What You'll Learn Here... Basics behind machine learning techniques Most commonly used machine learning algorithms, linear and logistic regression, decision trees support vector machines, k-nearest neighbors, random forests Solving multi-classification problems Data visualization with Matplotlib and data transformation with Pandas and Scikit-learn Solving multi-label classification problems And much, much more... Machine Learning With TensorFlow Here Is A Preview Of What You'll Learn Here... What is machine learning Main uses and benefits of machine learning How to get started with TensorFlow, installing and loading data Data flow graphs and basic TensorFlow expressions Creating MNIST classifiers with one-hot transformation And much, much more... Data Analytics With Python Here Is A Preview Of What You'll Learn Here... What is Data Analytics? Difference between data science, big data and data analytics Installing python Python data structures Pandas series and data frames And much, much more... Natural Language Processing With Python Here Is A Preview Of What You'll Learn Here... Challenges of natural language processing How natural language processing works? Part of speech tagging N-grams Running natural language processing script And much, much more... DevOps Handbook Here Is A Preview Of What You'll Learn Here... Issues and mistakes plaguing software development What is software development life cycle? How software development life cycle works? The origins of devops Testing and building systems tools And much, much more... DevOps Adoption Here Is A Preview Of What You'll Learn Here... Devops definition Overcoming traditional dev and ops Devops and security integration Devops success factors Is devops right for you? And much, much more... Get this book bundle NOW and SAVE money!

Imagine a world where you can make a computer program learn for itself? What if you were able to create any kind of program that you wanted, even as a beginner programmer, without all of the convoluted codes and other information that makes your head spin?

Build real-world applications with Python 2.7, CUDA 9, and CUDA 10. We suggest the use of Python 2.7 over Python 3.x, since Python 2.7 has stable support across all the libraries we use in this book. Key Features Expand your background in GPU programming—PyCUDA, scikit-cuda, and Nsight Effectively use CUDA libraries such as cuBLAS, cuFFT, and cuSolver Apply GPU programming to modern data science applications Book Description Hands-On GPU Programming with Python and CUDA hits the ground running: you'll start by learning how to apply Amdahl's Law, use a code profiler to identify bottlenecks in your Python code, and set up an appropriate GPU programming environment. You'll then see how to "query" the GPU's features and copy arrays of data to and from the GPU's own memory. As you make your way through the book, you'll launch code directly onto the GPU and write full blown GPU kernels and device functions in CUDA C. You'll get to grips with profiling GPU code effectively and fully test and debug your code using Nsight IDE. Next, you'll explore some of the more well-known NVIDIA libraries, such as cuFFT and cuBLAS. With a solid background in place, you will now apply your new-found knowledge to develop your very own GPU-based deep neural network from scratch. You'll then explore advanced topics, such as warp shuffling, dynamic parallelism, and PTX assembly. In the final chapter, you'll see some topics and applications related to GPU programming that you may wish to pursue, including AI, graphics, and blockchain. By the end of this book, you will be able to apply GPU programming to problems related to data science and high-performance computing. What you will learn Launch GPU code directly from Python Write effective and efficient GPU kernels and device functions Use libraries such as cuFFT, cuBLAS, and cuSolver Debug and profile your code with Nsight and Visual Profiler Apply GPU programming to datascience problems Build a GPU-based deep neuralnetwork from scratch Explore advanced GPU hardware features, such as warp shuffling Who this book is for Hands-On GPU Programming with Python and CUDA is for developers and data scientists who want to learn the basics of effective GPU programming to improve performance using Python code. You should have an understanding of first-year college or university-level engineering mathematics and physics, and have some experience with Python as well as in any C-based programming language such as C, C++, Go, or Java.

55% OFF for bookstores! Get maximum revenue from the sale of this book. Your customers will love this book.?

Do want to learn how machine learning and neural networks work quickly and simply? Do you want to know how to build a machine learning model and you have no programming skill? Do you know a bit of Python coding and want to learn more about how this deep learning works? This bundle is going to guide you to the basics and the principles behind machine learning. Machine learning is an active research domain and includes several different approaches. This bundle is going to help you understand the different approaches of machine learning and neural networks. It will guide you through the steps you need to build a machine learning model. This bundle is intended to address all these questions. In fact, the aim of this book is providing the absolute beginners or other programmers that has no experience with Python programming the basic and fundamental tools of the Python language. Through this bundle, you will learn: - Principles of machine learning - Types of machine learning: supervised, unsupervised, semi-supervised, and reinforcement learning - Advantages of each type of machine learning - Principle and types of neural networks - Steps to develop and fit artificial neural network model - Getting started and installing Python - Tools and platforms for Python programming - How to use pandas, NumPy and matplotlib Python libraries - How to develop a simple linear and logistic machine learning model - How to develop and train a multi-layer artificial neural network two ways: from scratch and using the Python libraries - When to use each type of machine learning - The general concept of artificial neural networks - Activation function in artificial neural network and to choose an activation function within an artificial neural network - The 5 main types of artificial neural network - The best function to be used to train artificial

neural networks. - the 2 main concepts to know in the training process of the artificial neural network - the main variants and algorithms for the formation of an artificial neural network and a machine learning model in general. - The basics of the three main Python languages that will help you get the work done-including TensorFlow, Keras, and PyTorch; - How to install the three Python libraries to help you get started; How to install and use magic command in Ipython Functionalities of NumPy library for numerical programming Functionalities of Pandas library for data analysis Even if you don't have any background in machine learning and Python programming, this book will give you the tools to develop machine learning models.

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.

Do you want to learn how machine learning and neural networks work quickly and simply? Do you want to know how to build a machine learning model, and you have no programming skills? Do you want to get started with learning data science? This book is going to guide you to the basics and the principles behind machine learning. Machine learning is an active research domain and includes several different approaches. This book is going to help you understand the various methods of machine learning and neural networks. It will guide you through the steps you need to build a machine learning model. Machine learning implies programming. This book will teach you Python programming. This book does not require any pre-programming skills. It will help to get you started in Python programming, as well as how to use Python libraries to analyze data and apply machine learning. Overall, this book is a go-to guide for getting started in machine learning modeling using Python programming. Once you get through the book, you will be able to develop your machine learning models using Python. Through this book, you will learn: - Principles of machine learning - Types of machine learning: supervised, unsupervised, semi-supervised, and reinforcement learning - Advantages of each type of machine learning - Principle and types of neural networks - Steps to develop and fit artificial neural network model - Getting started and installing Python - Tools and platforms for Python programming - How to use pandas, NumPy and matplotlib Python libraries - How to develop a simple linear and logistic machine learning model - How to build and train a multi-layer artificial neural network two ways: from scratch and using the Python libraries Even if you don't have any background in machine learning and Python programming, this book will give you the tools to develop machine learning models.

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-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.

55% OFF for bookstores! Get maximum revenue from the sale of this book. Your customers will love this book.

Programming With Python - 4 BOOK BUNDLE!! Deep Learning with Keras Here Is a Preview of What You'll Learn Here... The difference between deep learning and machine learning Deep neural networks Convolutional neural networks Building deep learning models with Keras Multi-layer perceptron network models Activation functions Handwritten recognition using MNIST Solving multi-class classification problems Recurrent neural networks and sequence classification And much more... Convolutional Neural Networks in Python Here Is a Preview of What You'll Learn In This Book... Convolutional neural networks structure How convolutional neural networks actually work Convolutional neural networks applications The importance of convolution operator Different convolutional neural networks layers and their importance Arrangement of spatial parameters How and when to use stride and zero-padding Method of parameter sharing Matrix multiplication and its importance Pooling and dense layers Introducing non-linearity relu activation function How to train your convolutional neural network models using backpropagation How and why to apply dropout CNN model training process How to build a convolutional neural network Generating predictions and calculating loss functions How to train and evaluate your MNIST classifier How to build a simple image classification CNN And much, much more! Python Machine Learning Here Is A Preview Of What You'll Learn Here... Basics behind machine learning techniques Different machine learning algorithms Fundamental machine learning applications and their importance Getting started with machine learning in Python, installing and starting SciPy Loading data and importing different libraries Data summarization and data visualization Evaluation of machine learning models and making predictions Most commonly used machine learning algorithms, linear and logistic regression, decision trees support vector machines, k-nearest neighbors, random forests Solving multi-clasifcation problems Data

visualization with Matplotlib and data transformation with Pandas and Scikit-learn Solving multi-label classification problems And much, much more... Machine Learning With TensorFlow Here Is a Preview of What You'll Learn Here... What is machine learning Main uses and benefits of machine learning How to get started with TensorFlow, installing and loading data Data flow graphs and basic TensorFlow expressions How to define your data flow graphs and how to use TensorBoard for data visualization Main TensorFlow operations and building tensors How to perform data transformation using different techniques How to build high performance data pipelines using TensorFlow Dataset framework How to create TensorFlow iterators Creating MNIST classifiers with one-hot transformation Get this book bundle NOW and SAVE money!

Create and unleash the power of neural networks by implementing C# and .Net code Key Features Get a strong foundation of neural networks with access to various machine learning and deep learning libraries Real-world case studies illustrating various neural network techniques and architectures used by practitioners Cutting-edge coverage of Deep Networks, optimization algorithms, convolutional networks, autoencoders and many more Book Description Neural networks have made a surprise comeback in the last few years and have brought tremendous innovation in the world of artificial intelligence. The goal of this book is to provide C# programmers with practical guidance in solving complex computational challenges using neural networks and C# libraries such as CNTK, and TensorFlowSharp. This book will take you on a step-by-step practical journey, covering everything from the mathematical and theoretical aspects of neural networks, to building your own deep neural networks into your applications with the C# and .NET frameworks. This book begins by giving you a quick refresher of neural networks. You will learn how to build a neural network from scratch using packages such as Encog, Aforge, and Accord. You will learn about various concepts and techniques, such as deep networks, perceptrons, optimization algorithms, convolutional networks, and autoencoders. You will learn ways to add intelligent features to your .NET apps, such as facial and motion detection, object detection and labeling, language understanding, knowledge, and intelligent search. Throughout this book, you will be working on interesting demonstrations that will make it easier to implement complex neural networks in your enterprise applications. What you will learn Understand perceptrons and how to implement them in C# Learn how to train and visualize a neural network using cognitive services Perform image recognition for detecting and labeling objects using C# and TensorFlowSharp Detect specific image characteristics such as a face using Accord.Net Demonstrate particle swarm optimization using a simple XOR problem and Encog Train convolutional neural networks using ConvNetSharp Find optimal parameters for your neural network functions using numeric and heuristic optimization techniques. Who this book is for This book is for Machine Learning Engineers, Data Scientists, Deep Learning Aspirants and Data Analysts who are now looking to move into advanced machine learning and deep learning with C#. Prior knowledge of machine learning and working experience with C# programming is required to take most out of this book

Build your Machine Learning portfolio by creating 6 cutting-edge Artificial Intelligence projects using neural networks in Python Key Features Discover neural network architectures (like CNN and LSTM) that are driving recent advancements in AI Build expert neural networks in Python using popular libraries such as Keras Includes projects such as object detection, face identification, sentiment analysis, and more Book Description Neural networks are at the core of recent AI advances, providing some of the best resolutions to many real-world problems, including image recognition, medical diagnosis, text analysis, and more. This book goes through some basic neural network and deep learning concepts, as well as some popular libraries in Python for implementing them. It contains practical demonstrations of neural networks in domains such as fare prediction, image classification, sentiment analysis, and more. In each case, the book provides a problem statement, the specific neural network architecture required to tackle that problem, the reasoning behind the algorithm used, and the associated Python code to implement the solution from scratch. In the process, you will gain hands-on experience with using popular Python libraries such as Keras to build and train your own neural networks from scratch. By the end of this book, you will have mastered the different neural network architectures and created cutting-edge AI projects in Python that will immediately strengthen your machine learning portfolio. What you will learn Learn various neural network architectures and its advancements in AI Master deep learning in Python by building and training neural network Master neural networks for regression and classification Discover convolutional neural networks for image recognition Learn sentiment analysis on textual data using Long Short-Term Memory Build and train a highly accurate facial recognition security system Who this book is for This book is a perfect match for data scientists, machine learning engineers, and deep learning enthusiasts who wish to create practical neural network projects in Python. Readers should already have some basic knowledge of machine learning and neural networks.

Create and unleash the power of neural networks by implementing professional, clean, and clear Java code About This Book* Learn to build amazing projects using neural networks including forecasting the weather and pattern recognition* Explore the Java multi-platform feature to run your personal neural networks everywhere* This step-by-step guide will help you solve real-world problems and links neural network theory to their application Who This Book Is For This book is for Java developers who want to know how to develop smarter applications using the power of neural networks. Those who deal with a lot of complex data and want to use it efficiently in their day-to-day apps will find this book quite useful. Some basic experience with statistical computations is expected. What You Will Learn* Develop an understanding of neural networks and how they can be fitted* Explore the learning process of neural networks* Build neural network applications with Java using hands-on examples* Discover the power of neural network's unsupervised learning process to extract the intrinsic knowledge hidden behind the data* Apply the code generated in practical examples, including weather forecasting and pattern recognition* Understand how to make the best choice of learning parameters to ensure you have a more effective application* Select and split data sets into training, test, and validation, and explore validation strategies In Detail Want to discover the current state-of-art in the field of neural networks that will let you understand and design new strategies to apply to more complex problems? This book takes you on a complete walkthrough of the process of developing basic to advanced practical examples based on neural networks with Java, giving you everything you need to stand out. You will first learn the basics of neural networks and their process of learning. We then focus on what Perceptrons are and their features. Next, you will implement self-organizing maps using practical examples. Further on, you will learn about some of the applications that are presented in this book such as weather forecasting, disease diagnosis, customer profiling, generalization, extreme machine learning, and characters recognition (OCR). Finally, you will learn methods to optimize and adapt neural networks in real time. All the examples generated in the book are provided in the form of illustrative source code, which merges object-oriented programming (OOP) concepts and neural network features to enhance your learning experience.

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

A step-by-step gentle journey through the mathematics of neural networks, and making your own using the Python computer language. Neural networks are a key element of deep learning and artificial intelligence, which today is capable of some truly impressive feats. Yet too few really understand how neural networks actually work. This guide will take you on a fun and unhurried journey, starting from very simple ideas, and gradually building up an understanding of how neural networks work. You won't need any mathematics beyond secondary school, and an accessible introduction to calculus is also included. The ambition of this guide is to make neural networks as accessible as possible to as many readers as possible - there are enough texts for advanced readers already! You'll learn to code in Python and make your own neural network, teaching it to recognise human handwritten numbers, and performing as well as professionally developed networks. Part 1 is about ideas. We introduce the mathematical ideas underlying the neural networks, gently with lots of illustrations and examples. Part 2 is practical. We introduce the popular and easy to learn Python programming language, and gradually builds up a neural network which can learn to recognise human handwritten numbers, easily getting it to perform as well as networks made by professionals. Part 3 extends these ideas further. We push the performance of our neural network to an industry leading 98% using only simple ideas and code, test the network on your own handwriting, take a privileged peek inside the mysterious mind of a neural network, and even get it all working on a Raspberry Pi. All the code in this has been tested to work on a Raspberry Pi Zero.

Become A Python Expert From Scratch! Python's popularity is growing tremendously and it's becoming more and more relevant economically and technologically. The fields of application of the language range from machine learning, over computer networking to business applications. In this 7 in 1 version you get a full collection of The Python Bible series. From the first volume on, you will be lead on a structured way to the mastery of Python. Besides the basics and the intermediate concepts, you will also learn how to apply it in areas like machine learning, financial analysis and neural networks. At the end you will additionally be introduced to one of the most interesting fields of computer science, which is computer vision After reading this collection, you will not only understand the programming language but you will also be able to work on projects in the stated fields. You will become a true Python expert! What You Will Learn: Beginner Level: - Basics of Programming with Python- Automation of Simple Processes- Programming of Modular Python Applications- Easy Transition to Other Languages (Java, C++ etc.) Intermediate Level: - Object-Oriented Programming- Network Programming- Penetration Testing with Python- Regular Expressions- Multithreading- XML Processing- Database Programming- Logging Data Science: - Analyzing and Processing Big Data- Statistical Calculations with Python- Visualization of Data- Working with NumPy, Matplotlib and Pandas Machine Learning: - Predicting Data with Machine Learning- Building Neural Networks with Tensorflow- Recognizing Handwritten Digits with Neural Networks- Applying Linear Models like Regression- K-Nearest-Neighbors Classification- K-Means Clustering- Support Vector Machines Finance: - Financial Analysis with Python- Analyzing and Graphing Stock Data- Plotting Trendlines- Predicting Share Prices with Machine Learning Neural Networks: - Generating Poetic d104s with Neural Networks- Predicting Sequential Data (Stocks, Weather etc.)- Processing Audio and Video Data- Recognizing Objects Like Horses, Cars and Trucks on Images- Understanding Recurrent Neural Networks- Understanding Convolutional Neural Networks Computer Vision: - Making unreadable texts readable again with thresholding- Extracting essential information out of images and videos- Edge detection- Template matching and feature matching- Movement detection in videos- Professional object recognition with OpenCV Start Your Journey And Become A Python Expert With The Python Bible!

Apply computer vision and machine learning concepts in developing business and industrial applications ?using a practical, step-by-step approach. The book comprises four main sections starting with setting up your programming environment and configuring your computer with all the prerequisites to run the code examples. Section 1 covers the basics of image and video processing with code examples of how to manipulate and extract useful information from the images. You will mainly use OpenCV with Python to work with examples in this section. Section 2

describes machine learning and neural network concepts as applied to computer vision. You will learn different algorithms of the neural network, such as convolutional neural network (CNN), region-based convolutional neural network (R-CNN), and YOLO. In this section, you will also learn how to train, tune, and manage neural networks for computer vision. Section 3 provides step-by-step examples of developing business and industrial applications, such as facial recognition in video surveillance and surface defect detection in manufacturing. The final section is about training neural networks involving a large number of images on cloud infrastructure, such as Amazon AWS, Google Cloud Platform, and Microsoft Azure. It walks you through the process of training distributed neural networks for computer vision on GPU-based cloud infrastructure. By the time you finish reading Building Computer Vision Applications Using Artificial Neural Networks and working through the code examples, you will have developed some real-world use cases of computer vision with deep learning. What You Will Learn · Employ image processing, manipulation, and feature extraction techniques · Work with various deep learning algorithms for computer vision · Train, manage, and tune hyperparameters of CNNs and object detection models, such as R-CNN, SSD, and YOLO · Build neural network models using Keras and TensorFlow · Discover best practices when implementing computer vision applications in business and industry · Train distributed models on GPU-based cloud infrastructure Who This Book Is For Data scientists, analysts, and machine learning and software engineering professionals with Python programming knowledge.

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

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.

Network Programming with Go teaches you how to write clean, secure network software with the programming language designed to make it seem easy. Go combines the best parts of many other programming languages. It's fast, scalable, and designed for high-performance networking and multiprocessing—in other words, it's perfect for network programming. Network Programming with Go is for developers ready to start leveraging Go's ease of use for writing secure, readable, production-ready network code. Early chapters establish a foundation of networking and traffic-routing know-how upon which the rest of the book builds. You'll put that knowledge to use as author Adam Woodbeck guides you through writing programs that communicate using TCP, UDP, Unix sockets, and other features that ensure reliable data transmission. As you progress, you'll explore higher-level network protocols like HTTP and HTTP/2, then build applications that securely interact with servers, clients, and APIs over a network using TLS. In addition, Woodbeck shows you how to create a simple messaging protocol, develop tools for monitoring network traffic, craft a custom web server, and implement best practices for interacting with cloud providers using their SDKs. Along the way, you'll learn:

- IP basics for writing effective network programs, such as IPv4 and IPv6 multicasting, ports, and network address translation
- How to use handlers, middleware, and multiplexers to build capable HTTP-based applications with minimal code
- The OSI and TCP/IP models for layered data architectures
- Methods for reading data from/writing data to a network connection, like the type-length-value encoding scheme
- Tools for incorporating authentication and encryption into your applications using TLS, like mutual authentication
- How to serialize data for storage or transmission in Go-friendly formats like JSON, Gob, XML, and protocol buffers
- How to Leverage Go's code generation support to efficiently communicate with gRPC-based network services

So get ready to take advantage of Go's built-in concurrency, rapid compiling, and rich standard library. Because when it comes to writing robust network programs, it's Go time.

This book is a guide on how to implement a neural network in the Python programming language. It begins by giving you a brief overview of neural networks so as to know what they are, where they are used, and how they are implemented. The next step is an exploration of the backpropagation algorithm. This is the algorithm behind the functionality of neural networks, and it involves a forward and backward pass. NumPy is a Python library which can be used for the purpose of implementation of a neural network. This library is discussed in this book, and you are guided on how to use it for that purpose. The functionality of neural networks has to be improved. The various ways to improve how a neural network works is also explored. You are then guided on how to implement neural networks with Neupy, another Python library. The following topics are discussed in this book: - A Brief Overview of Neural Networks - Backpropagation Algorithm - Neural Networks with Numpy - Improving a Neural Network in Python - Neupy - Models in Neural Networks

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

Build your Own Neural Network today. Through easy-to-follow instruction and examples, you'll learn the fundamentals of Deep learning and build your very own Neural Network in Python using TensorFlow, Keras, PyTorch, and Theano. While you have the option of spending thousands of dollars on big and boring textbooks, we recommend getting the same pieces of information for a fraction of the cost. So Get Your Copy Now!! Why this book? Book Objectives The following are the objectives of this book: To help you understand deep learning in detail To help you know how to get started with deep learning in Python by setting up the coding environment. To help you transition from a deep learning Beginner to a Professional. To help you learn how to develop a complete and functional artificial neural network model in Python on your own. Who this Book is for? The author targets the following groups of people: Anybody who is a complete beginner to deep learning with Python. Anybody in need of advancing their Python for deep learning skills. Professors, lecturers or tutors who are looking to find better ways to explain Deep Learning to their students in the simplest and easiest way. Students and academicians, especially those focusing on python programming, neural networks, machine learning, and deep learning. What do you need for this Book? You are required to have installed the following on your computer: Python 3.X. TensorFlow . Keras . PyTorch The Author guides you on how to install the rest of the Python libraries that are required for deep learning. The author will guide you on how to install and configure the rest. What is inside the book? What is Deep Learning? An Overview of Artificial Neural Networks. Exploring the Libraries. Installation and Setup. TensorFlow Basics. Deep Learning with TensorFlow. Keras Basics. PyTorch Basics. Creating Convolutional Neural Networks with PyTorch. Creating Recurrent Neural Networks with PyTorch. From the back cover. Deep learning is part of machine learning methods based on learning data representations. This book written by Samuel Burns provides an excellent introduction to deep learning methods for computer vision applications. The author does not focus on too much math since this guide is designed for developers who are beginners in the field of deep learning. The book has been grouped into chapters, with each chapter exploring a different feature of the deep learning libraries that can be used in Python programming language. Each chapter features a unique Neural Network architecture including Convolutional Neural Networks. After reading this book, you will be able to build your own Neural Networks using Tensorflow, Keras, and PyTorch. Moreover, the author has provided Python codes, each code performing a different task. Corresponding explanations have also been provided alongside each piece of code to help the reader understand the meaning of the various lines of the code. In addition to this, screenshots showing the output that each code should return have been given. The author has used a simple language to make it easy even for beginners to understand.

[Copyright: b75e48c5feaa431eb22c6f1b53bc40c9](https://www.amazon.com/dp/b75e48c5feaa431eb22c6f1b53bc40c9)