



model works, before sharing the code you need to implement them in PyTorch. This book is ideal if you want to rapidly add PyTorch to your deep learning toolset. What you will learn Use PyTorch to build: Simple Neural Networks – build neural networks the PyTorch way, with high-level functions, optimizers, and more Convolutional Neural Networks – create advanced computer vision systems Recurrent Neural Networks – work with sequential data such as natural language and audio Generative Adversarial Networks – create new content with models including SimpleGAN and CycleGAN Reinforcement Learning – develop systems that can solve complex problems such as driving or game playing Deep Learning workflows – move effectively from ideation to production with proper deep learning workflow using PyTorch and its utility packages Production-ready models – package your models for high-performance production environments Who this book is for Machine learning engineers who want to put PyTorch to work.

Apply modern deep learning techniques to build and train deep neural networks using Gorgonia Key Features Gain a practical understanding of deep learning using Golang Build complex neural network models using Go libraries and Gorgonia Take your deep learning model from design to deployment with this handy guide Book Description Go is an open source programming language designed by Google for handling large-scale projects efficiently. The Go ecosystem comprises some really powerful deep learning tools such as DQN and CUDA. With this book, you'll be able to use these tools to train and deploy scalable deep learning models from scratch. This deep learning book begins by introducing you to a variety of tools and libraries available in Go. It then takes you through building neural networks, including activation functions and the learning algorithms that make neural networks tick. In addition to this, you'll learn how to build advanced architectures such as autoencoders, restricted Boltzmann machines (RBMs), convolutional neural networks (CNNs), recurrent neural networks (RNNs), and more. You'll also understand how you can scale model deployments on the AWS cloud infrastructure for training and inference. By the end of this book, you'll have mastered the art of building, training, and deploying deep learning models in Go to solve real-world problems. What you will learn Explore the Go ecosystem of libraries and communities for deep learning Get to grips with Neural Networks, their history, and how they work Design and implement Deep Neural Networks in Go Get a strong foundation of concepts such as Backpropagation and Momentum Build Variational Autoencoders and Restricted Boltzmann Machines using Go Build models with CUDA and benchmark CPU and GPU models Who this book is for This book is for data scientists, machine learning engineers, and AI developers who want to build state-of-the-art deep learning models using Go. Familiarity with basic machine learning concepts and Go programming is required to get the best out of this book.

Hands-on Machine Learning with JavaScript

Build, Train, and Deploy End-to-end Machine Learning and Deep Learning Pipelines

Master deep learning algorithms with extensive math by implementing them using TensorFlow

Hands-On Machine Learning with R

Build intelligent computer vision applications using TensorFlow and Keras

A beginner's guide to building automated machine learning systems using AutoML and Python

Getting started with Microsoft ML.NET to implement popular machine learning algorithms in C#

**Implement supervised and unsupervised machine learning algorithms using C++ libraries such as PyTorch C++ API, Caffe2, Shogun, Shark-ML, mlpack, and dlib with the help of real-world examples and datasets** Key FeaturesBecome familiar with data processing, performance measuring, and model selection using various C++ librariesImplement practical machine learning and deep learning techniques to build smart modelsDeploy machine learning models to work on mobile and embedded devicesBook Description C++ can make your machine learning models run faster and more efficiently. This handy guide will help you learn the fundamentals of machine learning (ML), showing you how to use C++ libraries to get the most out of your data. This book makes machine learning with C++ for beginners easy with its example-based approach, demonstrating how to implement supervised and unsupervised ML algorithms through real-world examples. This book will get you hands-on with tuning and optimizing a model for different use cases, assisting you with model selection and the measurement of performance. You'll cover techniques such as product recommendations, ensemble learning, and anomaly detection using modern C++ libraries such as PyTorch C++ API, Caffe2, Shogun, Shark-ML, mlpack, and dlib. Next, you'll explore neural networks and deep learning using examples such as image classification and sentiment analysis, which will help you solve various problems. Later, you'll learn how to handle production and deployment challenges on mobile and cloud platforms, before discovering how to export and import models using the ONNX format. By the end of this C++ book, you will have real-world machine learning and C++ knowledge, as well as the skills to use C++ to build powerful ML systems. What you will learnExplore how to load and preprocess various data types to suitable C++ data structuresEmploy key machine learning algorithms with various C++ librariesUnderstand the grid-search approach to find the best parameters for a machine learning modelImplement an algorithm for filtering anomalies in user data using Gaussian distributionImprove collaborative filtering to deal with dynamic user preferencesUse C++ libraries and APIs to manage model structures and parametersImplement a C++ program to solve image classification tasks with LeNet architectureWho this book is for You will find this C++ machine learning book useful if you want to get started with machine learning algorithms and techniques using the popular C++ language. As well as being a useful first course in machine learning with C++, this book will also appeal to data analysts, data scientists, and machine learning developers who are looking to implement different machine learning models in production using varied datasets and examples. Working knowledge of the C++ programming language is mandatory to get started with this book.

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.

Create, train, and evaluate various machine learning models such as regression, classification, and clustering using ML.NET, Entity Framework, and ASP.NET Core Key FeaturesGet well-versed with the ML.NET framework and its components and APIs using practical examplesLearn how to build, train, and evaluate popular machine learning algorithms with ML.NET offeringsExtend your existing machine learning models by integrating with TensorFlow and other librariesBook Description Machine learning (ML) is widely used in many industries such as science, healthcare, and research and its popularity is only growing. In March 2018, Microsoft introduced ML.NET to help .NET enthusiasts in working with ML. With this book, you'll explore how to build ML.NET applications with the various ML models available using C# code. The book starts by giving you an overview of ML and the types of ML algorithms used, along with covering what ML.NET is and why you need it to build ML apps. You'll then explore the ML.NET framework, its components, and APIs. The book will serve as a practical guide to helping you build smart apps using the ML.NET library. You'll gradually become well versed in how to implement ML algorithms such as regression, classification, and clustering with real-world examples and datasets. Each chapter will cover the practical implementation, showing you how to implement ML within .NET applications. You'll also learn to integrate TensorFlow in ML.NET applications. Later you'll discover how to store the regression model housing price prediction result to the database and display the real-time predicted results from the database on your web application using ASP.NET Core Blazor and SignalR. By the end of this book, you'll have learned how to confidently perform basic to advanced-level machine learning tasks in ML.NET. What you will learnUnderstand the framework, components, and APIs of ML.NET using C#Develop regression models using ML.NET for employee attrition and file classificationEvaluate classification models for sentiment prediction of restaurant reviewsWork with clustering models for file type classificationsUse anomaly detection to find anomalies in both network traffic and login historyWork with ASP.NET Core Blazor to create an ML.NET enabled web applicationIntegrate pre-trained TensorFlow and ONNX models in a WPF ML.NET application for image classification and object detectionWho this book is for If you are a .NET developer who wants to implement machine learning models using ML.NET, then this book is for you. This book will also be beneficial for data scientists and machine learning developers who are looking for effective tools to implement various machine learning algorithms. A basic understanding of C# or .NET is mandatory to grasp the concepts covered in this book effectively.

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

Hands-On Transfer Learning with Python

Hands on Machine Learning with Python Tools, Concepts and Techniques

Build, train, and deploy end-to-end machine learning and deep learning pipelines

Hands-On Deep Learning for Images with TensorFlow

Build CNNs, RNNs, GANs, reinforcement learning, and more, quickly and easily

Hands-On Machine Learning on Google Cloud Platform

Data Science Fundamentals with Python

Aspiring data science professionals can learn the Scikit-Learn library along with the fundamentals of machine learning with this book. The book combines the Anaconda Python distribution with the popular Scikit-Learn library to demonstrate a wide range of supervised and unsupervised machine learning algorithms. Care is taken to walk you through the principles of machine learning through clear examples written in Python that you can try out and experiment with at home on your own machine. All applied math and programming skills required to master the content are covered in this book. In-depth knowledge of object-oriented programming is not required as working and complete examples are provided and explained. Coding examples are in-depth and complex when necessary. They are also concise, accurate, and complete, and complement the machine learning concepts introduced. Working the examples helps to build the skills necessary to understand and apply complex machine learning algorithms. Hands-on Scikit-Learn for Machine Learning Applications is an excellent starting point for those pursuing a career in machine learning. Students of this book will learn the fundamentals that are a prerequisite to competency. Readers will be exposed to the Anaconda distribution of Python that is designed specifically for data science professionals, and will build skills in the popular Scikit-Learn library that underlies many machine learning applications in the world of Python. What You'll LearnWork with simple and complex datasets common to Scikit-Learn Manipulate data into vectors and matrices for algorithmic processing Become familiar with the Anaconda distribution used in data scienceApply machine learning with Classifiers, Regressors, and Dimensionality Reduction Tune algorithms and find the best algorithms for each dataset Load data from and save to CSV, JSON, Numpy, and Pandas formats Who This Book Is For The aspiring data scientist yearning to break into machine learning through mastering the underlying fundamentals that are sometimes skipped over in the rush to be productive. Some knowledge of object-oriented programming and very basic applied linear algebra will make learning easier, although anyone can benefit from this book.

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. The updated edition of this best-selling book uses concrete examples, minimal theory, and two production-ready Python frameworks-Scikit-Learn and TensorFlow 2-to help you gain an intuitive understanding of the concepts and tools for building intelligent systems. Practitioners will learn a range of techniques that they can quickly put to use on the job. Part 1 employs Scikit-Learn to introduce fundamental machine learning tasks, such as simple linear regression. Part 2, which has been significantly updated, employs Keras and TensorFlow 2 to guide the reader through more advanced machine learning methods using 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. NEW FOR THE SECOND EDITION:Updated all code to TensorFlow 2 ; Introduced the high-level Keras API ; New and expanded coverage including TensorFlow's Data API, Eager Execution, Estimators API, deploying on Google Cloud ML, handling time series, embeddings and more With Early Release ebooks, you get books in their earliest form-the author's raw and unedited content as he or she writes-so you can take advantage of these technologies long before the official release of these titles. You'll also receive updates when significant changes are made, new chapters are available, and the final ebook bundle is released.

Deep learning simplified by taking supervised, unsupervised, and reinforcement learning to the next level using the Python ecosystem Key Features Build deep learning models with transfer learning principles in Python implement transfer learning to solve real-world research problems Perform complex operations such as image captioning neural style transfer Book Description Transfer learning is a machine learning (ML) technique where knowledge gained during training a set of problems can be used to solve other similar problems. The purpose of this book is two-fold; firstly, we focus on detailed coverage of deep learning (DL) and transfer learning, comparing and contrasting the two with easy-to-follow concepts and examples. The second area of focus is real-world examples and research problems using TensorFlow, Keras, and the Python ecosystem with hands-on examples. The book starts with the key essential concepts of ML and DL, followed by depiction and coverage of important DL architectures such as convolutional neural networks (CNNs), deep neural networks (DNNs), recurrent neural networks (RNNs), long short-term memory (LSTM), and capsule networks. Our focus then shifts to transfer learning concepts, such as model freezing, fine-tuning, pre-trained models including VGG, inception, ResNet, and how these systems perform better than DL models with practical examples. In the concluding chapters, we will focus on a multitude of real-world case studies and problems associated with areas such as computer vision, audio analysis and natural language processing (NLP). By the end of this book, you will be able to implement both DL and transfer learning principles in your own systems. What you will learn Set up your own DL environment with graphics processing unit (GPU) and Cloud support Delve into transfer learning principles with ML and DL models Explore various DL architectures, including CNN, LSTM, and capsule networks Learn about data and network representation and loss functions Get to grips with models and strategies in transfer learning Walk through potential challenges in building complex transfer learning models from scratch Explore real-world research problems related to computer vision and audio analysis Understand how transfer learning can be leveraged in NLP Who this book is for Hands-On Transfer Learning with Python is for data scientists, machine learning engineers, analysts and developers with an interest in data and applying state-of-the-art transfer learning methodologies to solve tough real-world problems. Basic proficiency in machine learning and Python is required.

Explore TensorFlow's capabilities to perform efficient deep learning on images Key Features Discover image processing for machine vision Build an effective image classification system using the power of CNNs Leverage TensorFlow's capabilities to perform efficient deep learning Book Description TensorFlow is Google's popular offering for machine learning and deep learning, quickly becoming a favorite tool for performing fast, efficient, and accurate deep learning tasks. Hands-On Deep Learning for Images with TensorFlow shows you the practical implementations of real-world projects, teaching you how to leverage TensorFlow's capabilities to perform efficient image processing using the power of deep learning. With the help of this book, you will get to grips with the different paradigms of performing deep learning such as deep neural nets and convolutional neural networks, followed by understanding how they can be implemented using TensorFlow. By the end of this book, you will have mastered all the concepts of deep learning and their implementation with TensorFlow and Keras. What you will learn Build machine learning models particularly focused on the MNIST digits Work with Docker and Keras to build an image classifier Understand natural language models to process text and images Prepare your dataset for machine learning Create classical, convolutional, and deep neural networks Create a RESTful image classification server Who this book is for Hands-On Deep Learning for Images with TensorFlow is for you if you are an application developer, data scientist, or machine learning practitioner looking to integrate machine learning into application software and master deep learning by implementing practical projects in TensorFlow. Knowledge of Python programming and basics of deep learning are required to get the best out of this book.

PyTorch Deep Learning Hands-On

Hands-On Deep Learning with Go

Hands-On Automated Machine Learning

Implement advanced deep learning and neural network models using TensorFlow and Keras

Hands-on Scikit-Learn for Machine Learning Applications

Build and deploy distributed deep learning applications on Apache Spark

**Explore reinforcement learning (RL) techniques to build cutting-edge games using Python libraries such as PyTorch, OpenAI Gym, and TensorFlow** Key FeaturesGet to grips with the different reinforcement and DRL algorithms for game developmentLearn how to implement components such as artificial agents, map and level generation, and audio generationGain insights into cutting-edge RL research and understand how it is similar to artificial general researchBook Description With the increased presence of AI in the gaming industry, developers are challenged to create highly responsive and adaptive games by integrating artificial intelligence into their projects. This book is your guide to learning how various reinforcement learning techniques and algorithms play an important role in game development with Python. Starting with the basics, this book will help you build a strong foundation in reinforcement learning for game development. Each chapter will assist you in implementing different reinforcement learning techniques, such as Markov decision processes (MDPs), Q-learning, actor-critic methods, SARSA, and deterministic policy gradient algorithms, to build logical self-learning agents. Learning these techniques will enhance your game development skills and add a variety of features to improve your game agent's productivity. As you advance, you'll understand how deep reinforcement learning (DRL) techniques can be used to devise strategies to help agents learn from their actions and build engaging games. By the end of this book, you'll be ready to apply reinforcement learning techniques to build a variety of projects and contribute to open source applications. What you will learnUnderstand how deep learning can be integrated into an RL agentExplore basic to advanced algorithms commonly used in game developmentBuild agents that can learn and solve problems in all types of environmentsTrain a Deep Q-Network (DQN) agent to solve the CartPole balancing problemDevelop game AI agents by understanding the mechanism behind complex AIIntegrate all the concepts learned into new projects or gaming agentsWho this book is for If you're a game developer looking to implement AI techniques to build next-generation games from scratch, this book is for you. Machine learning and deep learning practitioners, and RL researchers who want to understand how to use self-learning agents in the game domain will also find this book useful. Knowledge of game development and Python programming experience are required. Dig deep into the data with a hands-on guide to machine learning with updated examples and more! Machine Learning: Hands-On for Developers and Technical Professionals provides hands-on instruction and fully-coded working examples for the most common machine learning techniques used by developers and technical professionals. The book contains a breakdown of each ML variant, explaining how it works and how it is used within certain industries, allowing readers to incorporate the presented techniques into their own work as they follow along. A core tenant of machine learning is a strong focus on data preparation, and a full exploration of the various types of learning algorithms illustrates how the proper tools can help any developer extract information and insights from existing data. The book includes a full complement of Instructor's Materials to facilitate use in the classroom, making this resource useful for students and as a professional reference. At its core, machine learning is a mathematical, algorithm-based technology that forms the basis of historical data mining and modern big data science. Scientific analysis of big data requires a working knowledge of machine learning, which forms predictions based on known properties learned from training data. Machine Learning is an accessible, comprehensive guide for the non-mathematician, providing clear guidance that allows readers to: Learn the languages of machine learning including Hadoop, Mahout, and Weka Understand decision trees, Bayesian networks, and artificial neural networks Implement Association Rule, Real Time, and Batch learning Develop a strategic plan for safe, effective, and efficient machine learning By learning to construct a system that can learn from data, readers can increase their utility across industries. Machine learning sits at the core of deep dive data analysis and visualization, which is increasingly in demand as companies discover the goldmine hiding in their existing data. For the tech professional involved in data science, Machine Learning: Hands-On for Developers and Technical Professionals provides the skills and techniques required to dig deeper.

Reinforcement learning is a self-evolving type of machine learning that takes us closer to achieving true artificial intelligence. This easy-to-follow guide explains everything from scratch using rich examples written in Python.

Explore and implement deep learning to solve various real-world problems using modern R libraries such as TensorFlow, MXNet, H2O, and Deeptnet Key FeaturesUnderstand deep learning algorithms and architectures using R and determine which algorithm is best suited for a specific problemImprove models using parameter tuning, feature engineering, and ensemblingApply advanced neural network models such as deep autoencoders and generative adversarial networks (GANs) across different domainsBook Description Deep learning enables efficient and accurate learning from a massive amount of data. This book will help you overcome a number of challenges using various deep learning algorithms and architectures with R programming. This book starts with a brief overview of machine learning and deep learning and how to build your first neural network. You'll understand the architecture of various deep learning algorithms and their applicable fields, learn how to build deep learning models, optimize hyperparameters, and evaluate model performance. Various deep learning applications in image processing, natural language processing (NLP), recommendation systems, and predictive analytics will also be covered. Later chapters will show you how to tackle

recognition problems such as image recognition and signal detection, programmatically summarize documents, conduct topic modeling, and forecast stock market prices. Toward the end of the book, you will learn the common applications of GANs and how to build a face generation model using them. Finally, you'll get to grips with using reinforcement learning and deep reinforcement learning to solve various real-world problems. By the end of this deep learning book, you will be able to build and deploy your own deep learning applications using appropriate frameworks and algorithms. What you will learn

**Design a feedforward neural network to see how the activation function computes an output**

**Create an image recognition model using convolutional neural networks (CNNs)**

**Prepare data, decide hidden layers and neurons and train your model with the backpropagation algorithm**

**Apply text cleaning techniques to remove uninformative text using NLP**

**Build, train, and evaluate a GAN model for face generation**

**Understand the concept and implementation of reinforcement learning in R**

**Who this book is for** This book is for data scientists, machine learning engineers, and deep learning developers who are familiar with machine learning and are looking to enhance their knowledge of deep learning using practical examples. Anyone interested in increasing the efficiency of their machine learning applications and exploring various options in R will also find this book useful. Basic knowledge of machine learning techniques and working knowledge of the R programming language is expected.

**Building Skills for an AI-Driven Enterprise**

**Hands-On Reinforcement Learning with Python**

**Hands-On Unsupervised Learning Using Python**

**Create deep neural networks to solve computational problems using TensorFlow and Keras**

**Hands-On Deep Learning with Apache Spark**

**A practical guide to building and implementing neural network models using Go**

**Hands-On Machine Learning with Azure**

Traditional books on machine learning can be divided into two groups- those aimed at advanced undergraduates or early postgraduates with reasonable mathematical knowledge and those that are primers on how to code algorithms. The field is ready for a text that not only demonstrates how to use the algorithms that make up machine learning methods, but

This book introduces basic-to-advanced deep learning algorithms used in a production environment by AI researchers and principal data scientists; it explains algorithms intuitively, including the underlying math, and shows how to implement them using popular Python-based deep learning libraries such as TensorFlow.

While several market-leading companies have successfully transformed their business models by following data- and AI-driven paths, the vast majority have yet to reap the benefits. How can your business and analytics units gain a competitive advantage by capturing the full potential of this predictive revolution? This practical guide presents a battle-tested end-to-end method to help you translate business decisions into tractable prescriptive solutions using data and AI as fundamental inputs. Author Daniel Vaughan shows data scientists, analytics practitioners, and others interested in using AI to transform their businesses not only how to ask the right questions but also how to generate value using modern AI technologies and decision-making principles. You'll explore several use cases common to many enterprises, complete with examples you can apply when working to solve your own issues. Break business decisions into stages that can be tackled using different skills from the analytical toolbox Identify and embrace uncertainty in decision making and protect against common human biases Customize optimal decisions to different customers using predictive and prescriptive methods and technologies Ask business questions that create high value through AI- and data-driven technologies

Unleash Google's Cloud Platform to build, train and optimize machine learning models

**Key Features**

- Get well versed in GCP pre-existing services to build your own smart models
- A comprehensive guide covering aspects from data processing, analyzing to building and training ML models
- A practical approach to produce your trained ML models and port them to your mobile for easy access

**Book Description**

Google Cloud Machine Learning Engine combines the services of Google Cloud Platform with the power and flexibility of TensorFlow. With this book, you will not only learn to build and train different complexities of machine learning models at scale but also host them in the cloud to make predictions. This book is focused on making the most of the Google Machine Learning Platform for large datasets and complex problems. You will learn from scratch how to create powerful machine learning based applications for a wide variety of problems by leveraging different data services from the Google Cloud Platform. Applications include NLP, Speech to text, Reinforcement learning, Time series, recommender systems, image classification, video content inference and many other. We will implement a wide variety of deep learning use cases and also make extensive use of data related services comprising the Google Cloud Platform ecosystem such as Firebase, Storage APIs, Datalab and so forth. This will enable you to integrate Machine Learning and data processing features into your web and mobile applications. By the end of this book, you will know the main difficulties that you may encounter and get appropriate strategies to overcome these difficulties and build efficient systems. What you will learn

- Use Google Cloud Platform to build data-based applications for dashboards, web, and mobile
- Create, train and optimize deep learning models for various data science problems on big data
- Learn how to leverage BigQuery to explore big datasets
- Use Google's pre-trained TensorFlow models for NLP, image, video and much more
- Create models and architectures for Time series, Reinforcement Learning, and generative models
- Create, evaluate, and optimize TensorFlow and Keras models for a wide range of applications

**Who this book is for** This book is for data scientists, machine learning developers and AI developers who want to learn Google Cloud Platform services to build machine learning applications. Since the interaction with the Google ML platform is mostly done via the command line, the reader is supposed to have some familiarity with the bash shell and Python scripting. Some understanding of machine learning and data science concepts will be handy

**How Neural Networks Grow Smarter**

**Hands-On Machine Learning with ML.NET**

**Analytical Skills for AI and Data Science**

**A practical guide to designing, building, and improving neural network models using R**

**Hands-On Unsupervised Learning with Python**

**Hands-On Machine Learning with TensorFlow.js**

**A Guide for Data Scientists**

*Concepts, tools, and techniques to explore deep learning architectures and methodologies*

**Key Features**

- Explore advanced deep learning architectures using various datasets and frameworks
- Implement deep architectures for neural network models such as CNN, RNN, GAN, and many more
- Discover design patterns and different challenges for various deep learning architectures

**Book Description**

Deep learning architectures are composed of multilevel nonlinear operations that represent high-level abstractions; this allows you to learn useful feature representations from the data. This book will help you learn and implement deep learning architectures to resolve various deep learning research problems.

**Hands-On Deep Learning Architectures with Python** explains the essential learning algorithms used for deep and shallow architectures. Packed with practical implementations and ideas to help you build efficient artificial intelligence systems (AI), this book will help you learn how neural networks play a major role in building deep architectures. You will understand various deep learning architectures (such as AlexNet, VGG Net, GoogleNet) with easy-to-follow code and diagrams. In addition to this, the book will also guide you in building and training various deep architectures such as the Boltzmann mechanism, autoencoders, convolutional neural networks (CNNs), recurrent neural networks (RNNs), natural language processing (NLP), GAN, and more—all with practical implementations. By the end of this book, you will be able to construct deep models using popular frameworks and datasets with the required design patterns for each architecture. You will be ready to explore the potential of deep architectures in today's world. What you will learn

- Implement CNNs, RNNs, and other commonly used architectures with Python
- Explore architectures such as VGGNet, AlexNet, and GoogLeNet
- Build deep learning architectures for AI applications such as face and image recognition, fraud detection, and many more
- Understand the architectures and applications of Boltzmann machines and autoencoders with concrete examples
- Master artificial intelligence and neural network concepts and apply them to your architecture
- Understand deep learning architectures for mobile and embedded systems

**Who this book is for** If you're a data scientist, machine learning developer/engineer, or deep learning practitioner, or are curious about AI and want to upgrade your knowledge of various deep learning architectures, this book will appeal to you. You are expected to have some knowledge of statistics and machine learning algorithms to get the best out of this book

*Discover the skill-sets required to implement various approaches to Machine Learning with Python*

**Key Features**

- Explore unsupervised learning with clustering, autoencoders, restricted Boltzmann machines, and more
- Build your own neural network models using modern Python libraries

**Practical examples** show you how to implement different machine learning and deep learning techniques

**Book Description**

Unsupervised learning is about making use of raw, untagged data and applying learning algorithms to it to help a machine predict its outcome. With this book, you will explore the concept of unsupervised learning to cluster large sets of data and analyze them repeatedly until the desired outcome is found using Python. This book starts with the key differences between supervised, unsupervised, and semi-supervised learning. You will be introduced to the best-used libraries and frameworks from the Python ecosystem and address unsupervised learning in both the machine learning and deep learning domains. You will explore various algorithms, techniques that are used to implement unsupervised learning in real-world use cases. You will learn a variety of unsupervised learning approaches, including randomized optimization, clustering, feature selection and transformation, and information theory. You will get hands-on experience with how neural networks can be employed in unsupervised scenarios. You will also explore the steps involved in building and training a GAN in order to process images. By the end of this book, you will have learned the art of unsupervised learning for different real-world challenges. What you will learn

- Use cluster algorithms to identify and optimize natural groups of data
- Explore advanced non-linear and hierarchical clustering in action
- Soft label assignments for fuzzy c-means and Gaussian mixture models
- Detect anomalies through density estimation
- Perform principal component analysis using neural network models
- Create unsupervised models using GANs

**Who this book is for** This book is intended for statisticians, data scientists, machine learning developers, and deep learning practitioners who want to build smart applications by implementing key building block unsupervised learning, and master all the new techniques and algorithms offered in machine learning and deep learning using real-world examples. Some prior knowledge of machine learning concepts and statistics is desirable.

This practical guide will teach you how deep learning (DL) can be used to solve complex real-world problems. **Key Features**

- Explore deep reinforcement learning (RL), from the first principles to the latest algorithms
- Evaluate high-profile RL methods, including value iteration, deep Q-networks, policy gradients, TRPO, PPO, DDPG, D4PG, evolution strategies and genetic algorithms
- Keep up with the very latest industry developments, including AI-driven chatbots

**Book Description**

Recent developments in reinforcement learning (RL), combined with deep learning (DL), have seen unprecedented progress made towards training agents to solve complex problems in a human-like way. Google's use of algorithms to play and defeat the well-known Atari arcade games has propelled the field to prominence, and researchers are generating new ideas at a rapid pace. Deep Reinforcement Learning **Hands-On** is a comprehensive guide to the very latest DL tools and their limitations. You will evaluate methods including Cross-entropy and policy gradients, before applying them to real-world environments. Take on both the Atari set of virtual games and family favorites such as Connect4. The book provides an introduction to the basics of RL, giving you the know-how to code intelligent learning agents to take on a formidable array of practical tasks. Discover how to implement Q-learning on 'grid world' environments, teach your agent to buy and trade stocks, and find out how natural language models are driving the boom in chatbots. What you will learn

- Understand the DL context of RL and implement complex DL models
- Learn the foundation of RL: Markov decision processes
- Evaluate RL methods including Cross-entropy, DQN, Actor-Critic, TRPO, PPO, DDPG, D4PG and others
- Discover how to deal with discrete and continuous action spaces in various environments
- Defeat Atari arcade games using the value iteration method
- Create your own OpenAI Gym environment to train a stock trading agent
- Teach your agent to play Connect4 using AlphaGo Zero
- Explore the very latest deep RL research on topics including AI-driven chatbots

**Who this book is for** Some fluency in Python is assumed. Basic deep learning (DL) approaches should be familiar to readers and some practical experience in DL will be helpful. This book is an introduction to deep reinforcement learning (RL) and requires no background in RL.

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

**Concepts, Tools, and Techniques to Build Intelligent Systems**

O'Reilly Media, Inc."

**Hands-On for Developers and Technical Professionals**

**The Self-Assembling Brain**

**Hands-On Machine Learning for Cybersecurity**

**Concepts, Tools, and Techniques to Build Intelligent Systems**

**How to Build Applied Machine Learning Solutions from Unlabeled Data**

**Hands-On Machine Learning for Algorithmic Trading**

**Hands-On Machine Learning with C++**

**This book will help you explore how to implement different well-known machine learning algorithms with various C++ frameworks and libraries. You will cover basic to advanced machine learning concepts with practical and easy to follow examples. By the end of the book, you will be able to build various machine learning models with ease.**

**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.**

**"In this book, Peter Robin Hiesinger explores historical and contemporary attempts to understand the information needed to make biological and artificial neural networks. Developmental neurobiologists and computer scientists with an interest in artificial intelligence - driven by the promise and resources of biomedical research on the one hand, and by the promise and advances of computer technology on the other - are trying to understand the fundamental principles that guide the generation of an intelligent system. Yet, though researchers in these disciplines share a common interest, their perspectives and approaches are often quite different. The book makes the case that "the information problem" underlies both fields, driving the questions that are driving forward the frontiers, and aims to encourage cross-disciplinary communication and understanding, to help both fields make progress. The questions that challenge researchers in these fields include the following. How does genetic information unfold during the years-long process of human brain development, and can this be a short-cut to create human-level artificial intelligence? Is the biological brain just messy hardware that can be improved upon by running learning algorithms in computers? Can artificial intelligence bypass evolutionary programming of "grown" networks? These questions are tightly linked, and answering them requires an understanding of how information unfolds algorithmically to generate functional neural networks. Via a series of closely linked "discussions" (fictional dialogues between researchers in different disciplines) and pedagogical "seminars," the author explores the different challenges facing researchers working on neural networks, their different perspectives and approaches, as well as the common ground and understanding to be found amongst those sharing an interest in the development of biological brains and artificial intelligent systems"--**

**Integrate scikit-learn with various tools such as NumPy, pandas, imbalanced-learn, and scikit-surprise and use it to solve real-world machine learning problems**

**Key Features**

- Delve into machine learning with this comprehensive guide to scikit-learn and scientific Python
- Master the art of data-driven problem-solving with hands-on examples
- Foster your theoretical and practical knowledge of supervised and unsupervised machine learning algorithms

**Book Description**

Machine learning is applied everywhere, from business to research and academia, while scikit-learn is a versatile library that is popular among machine learning practitioners. This book serves as a practical guide for anyone looking to provide hands-on machine learning solutions with scikit-learn and Python toolkits. The book begins with an explanation of machine learning concepts and fundamentals, and strikes a balance between theoretical concepts and their applications. Each chapter covers a different set of algorithms, and shows you how to use them to solve real-life problems. You'll also learn about various key supervised and unsupervised machine learning algorithms using practical examples. Whether it is an instance-based learning algorithm, Bayesian estimation, a deep neural network, a tree-based ensemble, or a recommendation system, you'll gain a thorough understanding of its theory and learn when to apply it. As you advance, you'll learn how to deal with unlabeled data and when to use different clustering and anomaly detection algorithms. By the end of this machine learning book, you'll have learned how to take a data-driven approach to provide end-to-end machine learning solutions. You'll also have discovered how to formulate the problem at hand, prepare required data, and evaluate and deploy models in production. What you will learn

- Understand when to use supervised, unsupervised, or reinforcement learning algorithms
- Find out how to collect and prepare your data for machine learning tasks
- Tackle imbalanced data and optimize your algorithm for a bias or variance tradeoff
- Apply supervised and unsupervised algorithms to overcome various machine learning challenges
- Employ best practices for tuning your algorithm's hyper parameters
- Discover how to use neural networks for classification and regression
- Build, evaluate, and deploy your machine learning solutions to production

**Who this book is for** This book is for data scientists, machine learning practitioners, and anyone who wants to learn how machine learning algorithms work and to build different machine learning models using the Python ecosystem. The book will help you take your knowledge of machine learning to the next level by grasping its ins and outs and tailoring it to your needs.

**Working knowledge of Python and a basic understanding of underlying mathematical and statistical concepts is required.**

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

**Hands-On Deep Learning Algorithms with Python**

**Hands-On Deep Learning Architectures with Python**

**Build powerful models with cognitive machine learning and artificial intelligence**

**Implementing self-learning agents in games using artificial intelligence techniques**

**Hands-On Machine Learning with Scikit-Learn and TensorFlow**