
Hands On Machine Learning With Scikit Learn And Tensorflow Concepts Tools And Techniques For Building Intelligent Systems

As recognized, adventure as well as experience just about lesson, amusement, as without difficulty as harmony can be gotten by just checking out a books **Hands On Machine Learning With Scikit Learn And Tensorflow Concepts Tools And Techniques For Building Intelligent Systems** next it is not directly done, you could agree to even more going on for this life, concerning the world.

We have the funds for you this proper as well as simple exaggeration to get those all. We present Hands On Machine Learning With Scikit Learn And Tensorflow Concepts Tools And Techniques For Building Intelligent Systems and numerous books collections from fictions to scientific research in any way. in the middle of them is this Hands On Machine Learning With Scikit Learn And Tensorflow Concepts Tools And Techniques For Building Intelligent Systems that can be your partner.

*Hands On Machine Learning With
Scikit Learn And Tensorflow Concepts
Tools And Techniques For Building
Intelligent Systems*

2022-05-07

WALKER GAMBLE

Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow, 2nd Edition John Wiley & Sons

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
Machine Learning O'Reilly Media

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.

Packt Publishing Ltd

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
Hands-On Reinforcement Learning for Games Packt Publishing Ltd

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 Learn Work 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 science Apply 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.

Analytical Skills for AI and Data Science CRC Press

Deep learning is often viewed as the exclusive domain of math PhDs and big tech companies. But as this hands-on guide

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

[Grokking Deep Learning](#) Packt Publishing Ltd

Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow "O'Reilly Media, Inc."

Hands-On Deep Learning with Apache Spark "O'Reilly Media, Inc."

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.

[Hands-On Machine Learning with Scikit-Learn, Keras, and](#)

[TensorFlow](#) "O'Reilly Media, Inc."

Through a recent series of 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 bestselling book uses concrete examples, minimal theory, and production-ready Python frameworks (Scikit-Learn, Keras, and TensorFlow) to help you gain an intuitive understanding of the concepts and tools for building intelligent systems. With this updated third edition, author Aurélien Géron explores a range of techniques, starting with simple linear regression and progressing to deep neural networks. Numerous code examples and exercises throughout the book help you apply what you've learned. Programming experience is all you need to get started. Use Scikit-learn to track an example ML project end to end. Explore several models, including support vector machines, decision trees, random forests, and ensemble methods. Exploit unsupervised learning techniques such as dimensionality reduction, clustering, and anomaly detection. Dive into neural net architectures, including convolutional nets, recurrent nets, generative adversarial networks, autoencoders, diffusion models, and transformers. Use TensorFlow and Keras to build and train neural nets for computer vision, natural language processing, generative models, and deep reinforcement learning.

[Hands-On Machine Learning with Microsoft Excel 2019](#) Packt Publishing Ltd

Get into the world of smart data security using machine learning algorithms and Python libraries. Key Features: Learn machine

learning algorithms and cybersecurity fundamentals Automate your daily workflow by applying use cases to many facets of security Implement smart machine learning solutions to detect various cybersecurity problems Book Description Cyber threats today are one of the costliest losses that an organization can face. In this book, we use the most efficient tool to solve the big problems that exist in the cybersecurity domain. The book begins by giving you the basics of ML in cybersecurity using Python and its libraries. You will explore various ML domains (such as time series analysis and ensemble modeling) to get your foundations right. You will implement various examples such as building system to identify malicious URLs, and building a program to detect fraudulent emails and spam. Later, you will learn how to make effective use of K-means algorithm to develop a solution to detect and alert you to any malicious activity in the network. Also learn how to implement biometrics and fingerprint to validate whether the user is a legitimate user or not. Finally, you will see how we change the game with TensorFlow and learn how deep learning is effective for creating models and training systems What you will learn Use machine learning algorithms with complex datasets to implement cybersecurity concepts Implement machine learning algorithms such as clustering, k-means, and Naive Bayes to solve real-world problems Learn to speed up a system using Python libraries with NumPy, Scikit-learn, and CUDA Understand how to combat malware, detect spam, and fight financial fraud to mitigate cyber crimes Use TensorFlow in the cybersecurity domain and implement real-world examples Learn how machine learning and Python can be used in complex cyber issues Who this book is for This book is for the data scientists,

machine learning developers, security researchers, and anyone keen to apply machine learning to up-skill computer security. Having some working knowledge of Python and being familiar with the basics of machine learning and cybersecurity fundamentals will help to get the most out of the book [Hands-On Machine Learning for Cybersecurity](#) Packt Publishing Ltd

Implement machine learning, cognitive services, and artificial intelligence solutions by leveraging Azure cloud technologies Key Features Learn advanced concepts in Azure ML and the Cortana Intelligence Suite architecture Explore ML Server using SQL Server and HDInsight capabilities Implement various tools in Azure to build and deploy machine learning models Book Description Implementing Machine learning (ML) and Artificial Intelligence (AI) in the cloud had not been possible earlier due to the lack of processing power and storage. However, Azure has created ML and AI services that are easy to implement in the cloud. Hands-On Machine Learning with Azure teaches you how to perform advanced ML projects in the cloud in a cost-effective way. The book begins by covering the benefits of ML and AI in the cloud. You will then explore Microsoft's Team Data Science Process to establish a repeatable process for successful AI development and implementation. You will also gain an understanding of AI technologies available in Azure and the Cognitive Services APIs to integrate them into bot applications. This book lets you explore prebuilt templates with Azure Machine Learning Studio and build a model using canned algorithms that can be deployed as web services. The book then takes you through a preconfigured series of virtual machines in Azure targeted at AI development

scenarios. You will get to grips with the ML Server and its capabilities in SQL and HDInsight. In the concluding chapters, you'll integrate patterns with other non-AI services in Azure. By the end of this book, you will be fully equipped to implement smart cognitive actions in your models. What you will learn

Discover the benefits of leveraging the cloud for ML and AI
Use Cognitive Services APIs to build intelligent bots
Build a model using canned algorithms from Microsoft and deploy it as a web service
Deploy virtual machines in AI development scenarios
Apply R, Python, SQL Server, and Spark in Azure
Build and deploy deep learning solutions with CNTK, MMLSpark, and TensorFlow
Implement model retraining in IoT, Streaming, and Blockchain solutions
Explore best practices for integrating ML and AI functions with ADLA and logic apps

Who this book is for
If you are a data scientist or developer familiar with Azure ML and cognitive services and want to create smart models and make sense of data in the cloud, this book is for you. You'll also find this book useful if you want to bring powerful machine learning services into your cloud applications. Some experience with data manipulation and processing, using languages like SQL, Python, and R, will aid in understanding the concepts covered in this book

Hands-On Machine Learning with C++ Packt Publishing Ltd
Automate data and model pipelines for faster machine learning applications
Key Features
Build automated modules for different machine learning components
Understand each component of a machine learning pipeline in depth
Learn to use different open source AutoML and feature engineering platforms

Book Description
AutoML is designed to automate parts of Machine Learning. Readily available AutoML tools are making data science

practitioners' work easy and are received well in the advanced analytics community. Automated Machine Learning covers the necessary foundation needed to create automated machine learning modules and helps you get up to speed with them in the most practical way possible. In this book, you'll learn how to automate different tasks in the machine learning pipeline such as data preprocessing, feature selection, model training, model optimization, and much more. In addition to this, it demonstrates how you can use the available automation libraries, such as auto-sklearn and MLBox, and create and extend your own custom AutoML components for Machine Learning. By the end of this book, you will have a clearer understanding of the different aspects of automated Machine Learning, and you'll be able to incorporate automation tasks using practical datasets. You can leverage your learning from this book to implement Machine Learning in your projects and get a step closer to winning various machine learning competitions. What you will learn

Understand the fundamentals of Automated Machine Learning systems
Explore auto-sklearn and MLBox for AutoML tasks
Automate your preprocessing methods along with feature transformation
Enhance feature selection and generation using the Python stack
Assemble individual components of ML into a complete AutoML framework
Demystify hyperparameter tuning to optimize your ML models
Dive into Machine Learning concepts such as neural networks and autoencoders
Understand the information costs and trade-offs associated with AutoML

Who this book is for
If you're a budding data scientist, data analyst, or Machine Learning enthusiast and are new to the concept of automated machine learning, this book is ideal for you. You'll also find this

book useful if you're an ML engineer or data professional interested in developing quick machine learning pipelines for your projects. Prior exposure to Python programming will help you get the best out of this book.

[Hands-On Deep Learning Architectures with Python](#) Packt Publishing Ltd

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.

Hands-on Scikit-Learn for Machine Learning Applications Packt Publishing Ltd

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
Deep Reinforcement Learning Hands-On Packt Publishing Ltd
"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"--

Hands-On Deep Learning for Images with TensorFlow CRC Press
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.

Deep Learning for Coders with fastai and PyTorch Packt Publishing Ltd

This book covers the fundamentals of machine learning with Python in a concise and dynamic manner. It covers data mining and large-scale machine learning using Apache Spark. About This Book Take your first steps in the world of data science by understanding the tools and techniques of data analysis Train efficient Machine Learning models in Python using the supervised and unsupervised learning methods Learn how to use Apache Spark for processing Big Data efficiently Who This Book Is For If you are a budding data scientist or a data analyst who wants to analyze and gain actionable insights from data using Python, this book is for you. Programmers with some experience in Python who want to enter the lucrative world of Data Science will also find this book to be very useful, but you don't need to be an expert Python coder or mathematician to get the most from this book. What You Will Learn Learn how to clean your data and ready it for analysis Implement the popular clustering and regression methods in Python Train efficient machine learning models using decision trees and random forests Visualize the results of your analysis using Python's Matplotlib library Use Apache Spark's MLlib package to perform machine learning on large datasets In Detail Join Frank Kane, who worked on Amazon and IMDb's machine learning algorithms, as he guides you on your first steps into the world of data science. Hands-On Data Science and Python Machine Learning gives you the tools that you need to understand and explore the core topics in the field, and the confidence and practice to build and analyze your own machine learning models. With the help of interesting and easy-

to-follow practical examples, Frank Kane explains potentially complex topics such as Bayesian methods and K-means clustering in a way that anybody can understand them. Based on Frank's successful data science course, Hands-On Data Science and Python Machine Learning empowers you to conduct data analysis and perform efficient machine learning using Python. Let Frank help you unearth the value in your data using the various data mining and data analysis techniques available in Python, and to develop efficient predictive models to predict future results. You will also learn how to perform large-scale machine learning on Big Data using Apache Spark. The book covers preparing your data for analysis, training machine learning models, and visualizing the final data analysis. Style and approach This comprehensive book is a perfect blend of theory and hands-on code examples in Python which can be used for your reference at any time.

[Hands-On Machine Learning with TensorFlow.js](#) Packt Publishing Ltd

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.

[Hands-On Deep Learning with TensorFlow](#) Princeton University Press

Many industry experts consider unsupervised learning the next frontier in artificial intelligence, one that may hold the key to general artificial intelligence. Since the majority of the world's data is unlabeled, conventional supervised learning cannot be applied. Unsupervised learning, on the other hand, can be applied to unlabeled datasets to discover meaningful patterns buried deep in the data, patterns that may be near impossible for humans to uncover. Author Ankur Patel shows you how to apply unsupervised learning using two simple, production-ready Python frameworks: Scikit-learn and TensorFlow using Keras. With code and hands-on examples, data scientists will identify difficult-to-find patterns in data and gain deeper business insight, detect anomalies, perform automatic feature engineering and selection, and generate synthetic datasets. All you need is programming and some machine learning experience to get started. Compare the strengths and weaknesses of the different machine learning approaches: supervised, unsupervised, and reinforcement learning Set up and manage machine learning projects end-to-end Build an anomaly detection system to catch credit card fraud Clusters users into distinct and homogeneous groups Perform semisupervised learning Develop movie recommender systems using restricted Boltzmann machines Generate synthetic images using generative adversarial networks

Hands-On Deep Learning with Go Packt Publishing Ltd

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

The Self-Assembling Brain Simon and Schuster

The ability to crunch data effectively can propel your career or business to great heights. Machine Learning is the most effective data analysis tool. While it is a complex topic, it can be broken down into simpler steps, as show in this book. We are using Python, which is a great programming language for beginners.