Machine Learning with Python Cookbook: Recipes for Real-World Applications

Machine learning is a powerful tool that allows computers to learn from data without being explicitly programmed. This makes it ideal for a wide range of tasks, from predicting customer churn to detecting fraud.

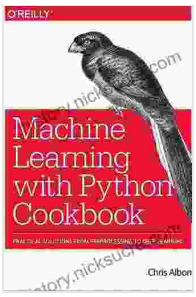

Machine Learning with Python Cookbook: Practical Solutions from Preprocessing to Deep Learning

by Chris Albon

4.6 out of 5

Language	: English
File size	: 4797 KB
Text-to-Speech	: Enabled
Screen Reader	: Supported
Enhanced typesetting	: Enabled
Print length	: 368 pages

Python is a popular programming language for machine learning because it is easy to learn and use, and there are a wide range of libraries and frameworks available for machine learning tasks.

This cookbook provides a collection of recipes that will help you master machine learning with Python. The recipes cover a wide range of topics, from basic data preprocessing to advanced deep learning techniques.

What You'll Learn

This cookbook will teach you how to:

* Preprocess data for machine learning * Train and evaluate machine learning models * Use machine learning models to make predictions * Deploy machine learning models to production

Who This Book Is For

This book is for anyone who wants to learn how to use machine learning with Python. Whether you're a beginner or an experienced data scientist, you'll find something in this book to help you.

Table of Contents

Chapter 1: Data Preprocessing

The first step in any machine learning project is to preprocess the data. This involves cleaning the data, removing outliers, and transforming the data into a format that is suitable for machine learning algorithms.

This chapter provides recipes for the following data preprocessing tasks:

* Cleaning data * Removing outliers * Transforming data * Scaling data * Normalizing data

Chapter 2: Linear Regression

Linear regression is a simple but powerful machine learning algorithm that can be used to predict continuous values. This chapter provides recipes for

the following linear regression tasks:

* Fitting a linear regression model * Evaluating a linear regression model * Using a linear regression model to make predictions

Chapter 3: Logistic Regression

Logistic regression is a machine learning algorithm that can be used to predict binary outcomes. This chapter provides recipes for the following logistic regression tasks:

* Fitting a logistic regression model * Evaluating a logistic regression model * Using a logistic regression model to make predictions

Chapter 4: Decision Trees

Decision trees are a machine learning algorithm that can be used to predict both continuous and binary outcomes. This chapter provides recipes for the following decision tree tasks:

* Fitting a decision tree model * Evaluating a decision tree model * Using a decision tree model to make predictions

Chapter 5: Random Forests

Random forests are a machine learning algorithm that combines multiple decision trees to improve accuracy. This chapter provides recipes for the following random forest tasks:

* Fitting a random forest model * Evaluating a random forest model * Using a random forest model to make predictions

Chapter 6: Support Vector Machines

Support vector machines are a machine learning algorithm that can be used to predict both continuous and binary outcomes. This chapter provides recipes for the following support vector machine tasks:

* Fitting a support vector machine model * Evaluating a support vector machine model * Using a support vector machine model to make predictions

Chapter 7: Naive Bayes

Naive Bayes is a machine learning algorithm that can be used to predict both continuous and binary outcomes. This chapter provides recipes for the following Naive Bayes tasks:

* Fitting a Naive Bayes model * Evaluating a Naive Bayes model * Using a Naive Bayes model to make predictions

Chapter 8: Clustering

Clustering is a machine learning technique that can be used to group data points into similar groups. This chapter provides recipes for the following clustering tasks:

* K-means clustering * Hierarchical clustering * Density-based spatial clustering

Chapter 9: Dimensionality Reduction

Dimensionality reduction is a machine learning technique that can be used to reduce the number of features in a dataset. This chapter provides recipes for the following dimensionality reduction tasks:

* Principal component analysis * Linear discriminant analysis * Singular value decomposition

Chapter 10: Deep Learning

Deep learning is a machine learning technique that uses artificial neural networks to learn from data. This chapter provides recipes for the following deep learning tasks:

* Convolutional neural networks * Recurrent neural networks * Autoencoders

This cookbook provides a comprehensive collection of recipes that will help you master machine learning with Python. The recipes cover a wide range of topics, from basic data preprocessing to advanced deep learning techniques.

Whether you're a beginner or an experienced data scientist, you'll find something in this book to help you.

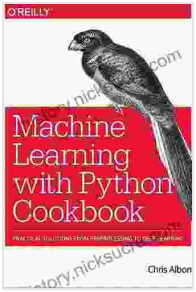

Machine Learning with Python Cookbook: Practical Solutions from Preprocessing to Deep Learning

by Chris Albon

4.6 out of 5

Language	: English
File size	: 4797 KB
Text-to-Speech	: Enabled
Screen Reader	: Supported
Enhanced typesetting	: Enabled
Print length	: 368 pages

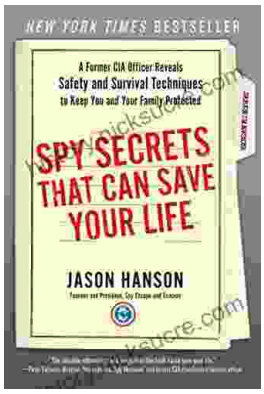

Spy Secrets That Can Save Your Life

` In the world of espionage, survival is paramount. Intelligence operatives face life-threatening situations on a regular basis, and they rely...

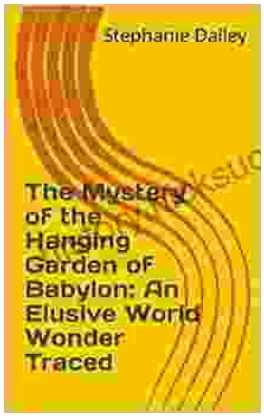

An Elusive World Wonder Traced

For centuries, the Hanging Gardens of Babylon have been shrouded in mystery. Now, researchers believe they have finally pinpointed the location of...