

Studio Notebook User Guide

Recognizing the showing off ways to get this books **Studio Notebook User Guide** is additionally useful. You have remained in right site to begin getting this info. acquire the Studio Notebook User Guide link that we have the funds for here and check out the link.

You could purchase guide Studio Notebook User Guide or acquire it as soon as feasible. You could quickly download this Studio Notebook User Guide after getting deal. So, later than you require the books swiftly, you can straight get it. Its appropriately unconditionally simple and consequently fats, isnt it? You have to favor to in this tone



[fastText Quick Start Guide](#) Packt Publishing Ltd

Automate the process of building, training, and deploying machine learning applications to production with AWS solutions such as SageMaker Autopilot, AutoGluon, Step Functions, Amazon Managed Workflows for Apache Airflow, and more Key FeaturesExplore the various AWS services that make automated machine learning easierRecognize the role of DevOps and MLOps methodologies in pipeline automationGet acquainted with additional AWS services such as Step Functions, MWAA, and more to overcome automation challengesBook Description AWS provides a wide range of solutions to help automate a machine learning workflow with just a few lines of code. With this practical book, you'll learn how to automate a machine learning pipeline using the various AWS services. Automated Machine Learning on AWS begins with a quick overview of what the machine learning pipeline/process looks like and highlights the typical challenges that you may face when building a pipeline. Throughout the book, you'll become well versed with various AWS solutions such as Amazon SageMaker Autopilot, AutoGluon, and AWS Step Functions to automate an end-to-end ML process with the help of hands-on examples. The book will show you how to build, monitor, and execute a CI/CD pipeline for the ML process and how the various CI/CD services within AWS can be applied to a use case with the Cloud Development Kit (CDK). You'll understand what a data-centric ML process is by working with the Amazon Managed Services for Apache Airflow and then build a managed Airflow environment. You'll also cover the key success criteria for an MLSDLC implementation and the process of creating a self-mutating CI/CD pipeline using AWS CDK from the perspective of the platform engineering team. By the end of this AWS book, you'll be able to effectively automate a complete machine learning pipeline and deploy it to production. What you will learnEmploy SageMaker Autopilot and Amazon SageMaker SDK to automate the machine learning processUnderstand how to use AutoGluon to automate complicated model building tasksUse the AWS CDK to codify the machine learning processCreate, deploy, and rebuild a CI/CD pipeline on AWSBuild an ML workflow using AWS Step Functions and the Data Science SDKLeverage the Amazon SageMaker Feature Store to automate the machine learning software development life cycle (MLSDLC)Discover how to use Amazon MWAA for a data-centric ML processWho this book is for This book is for the novice as well as experienced machine learning practitioners looking to automate the process of building, training, and deploying machine learning-based solutions into production, using both purpose-built and other AWS services. A basic understanding of the end-to-end machine learning process and concepts, Python programming, and AWS is necessary to make the most out of this book.

[Julia 1.0 Programming Complete Reference Guide](#) "O'Reilly Media, Inc."

Enhance your SQL query writing skills to provide greater business value using advanced techniques such as common table expressions, window functions, and JSON Purchase of the print or Kindle book includes a free PDF eBook Key Features Examine query design and performance using query plans and indexes Solve business problems using advanced techniques such as common table expressions and window functions Use SQL in modern data platform solutions with JSON and Jupyter notebooks Book Description SQL has been the de facto standard when interacting with databases for decades and shows no signs of going away. Through the years, report developers or data wranglers have had to learn SQL on the fly to meet the business needs, so if you are someone who needs to write queries, SQL Query Design and Pattern Best Practices is for you. This book will guide you through making efficient SQL queries by reducing set sizes for effective results. You'll learn how to format your results to make them easier to consume at their destination. From there, the book will take you through solving complex business problems using more advanced techniques, such as common table expressions and window functions, and advance to uncovering issues resulting from security in the underlying dataset. Armed with this knowledge, you'll have a foundation for building queries and be ready to shift focus to using tools, such as query plans and indexes, to optimize those queries. The book will go over the modern data estate, which includes data lakes and JSON data, and wrap up with a brief on how to use Jupyter notebooks in your SQL journey. By the end of this SQL book, you'll be able to make efficient SQL queries that will improve your report writing and the overall

SQL experience. What you will learn Build efficient queries by reducing the data being returned Manipulate your data and format it for easier consumption Form common table expressions and window functions to solve complex business issues Understand the impact of SQL security on your results Understand and use query plans to optimize your queries Understand the impact of indexes on your query performance and design Work with data lake data and JSON in SQL queries Organize your queries using Jupyter notebooks Who this book is for This book is for SQL developers, data analysts, report writers, data scientists, and other data gatherers looking to expand their skills for complex querying as well as for building more efficient and performant queries. For those new to SQL, this book can help you accelerate your learning and keep you from making common mistakes.

A Reference Guide to Television's Bonanza "O'Reilly Media, Inc."

Build an end-to-end geospatial data lake in AWS using popular AWS services such as RDS, Redshift, DynamoDB, and Athena to manage geodata Purchase of the print or Kindle book includes a free PDF eBook. Key Features Explore the architecture and different use cases to build and manage geospatial data lakes in AWS Discover how to leverage AWS purpose-built databases to store and analyze geospatial data Learn how to recognize which anti-patterns to avoid when managing geospatial data in the cloud Book DescriptionManaging geospatial data and building location-based applications in the cloud can be a daunting task. This comprehensive guide helps you overcome this challenge by presenting the concept of working with geospatial data in the cloud in an easy-to-understand way, along with teaching you how to design and build data lake architecture in AWS for geospatial data. You ' ll begin by exploring the use of AWS databases like Redshift and Aurora PostgreSQL for storing and analyzing geospatial data. Next, you ' ll leverage services such as DynamoDB and Athena, which offer powerful built-in geospatial functions for indexing and querying geospatial data. The book is filled with practical examples to illustrate the benefits of managing geospatial data in the cloud. As you advance, you ' ll discover how to analyze and visualize data using Python and R, and utilize QuickSight to share derived insights. The concluding chapters explore the integration of commonly used platforms like Open Data on AWS, OpenStreetMap, and ArcGIS with AWS to enable you to optimize efficiency and provide a supportive community for continuous learning. By the end of this book, you ' ll have the necessary tools and expertise to build and manage your own geospatial data lake on AWS, along with the knowledge needed to tackle geospatial data management challenges and make the most of AWS services.What you will learn Discover how to optimize the cloud to store your geospatial data Explore management strategies for your data repository using AWS Single Sign-On and IAM Create effective SQL queries against your geospatial data using Athena Validate postal addresses using Amazon Location services Process structured and unstructured geospatial data efficiently using R Use Amazon SageMaker to enable machine learning features in your application Explore the free and subscription satellite imagery data available for use in your GIS Who this book is forIf you understand the importance of accurate coordinates, but not necessarily the cloud, then this book is for you. This book is best suited for GIS developers, GIS analysts, data analysts, and data scientists looking to enhance their solutions with geospatial data for cloud-centric applications. A basic understanding of geographic concepts is suggested, but no experience with the cloud is necessary for understanding the concepts in this book.

[The Practitioner's Guide to Graph Data](#) "O'Reilly Media, Inc."

bookdown: Authoring Books and Technical Documents with R Markdown presents a much easier way to write books and technical publications than traditional tools such as LaTeX and Word. The bookdown package inherits the simplicity of syntax and flexibility for data analysis from R Markdown, and extends R Markdown for technical writing, so that you can make better use of document elements such as figures, tables, equations, theorems, citations, and references. Similar to LaTeX, you can number and cross-reference these elements with bookdown. Your document can even include live examples so readers can interact with them while reading the book. The book can be rendered to multiple output formats, including LaTeX/PDF, HTML, EPUB, and Word, thus making it easy to put your documents online. The style and theme of these output formats can be customized. We used books and R primarily for examples in this book, but bookdown is not only for books or R. Most features introduced in this book also apply to other types of publications: journal papers, reports, dissertations, course handouts, study notes, and even novels. You do not have to use R, either. Other choices of computing languages include Python, C, C++, SQL, Bash, Stan, JavaScript, and so on, although R is best supported. You can also leave out computing, for example, to write a fiction. This book itself is an example of publishing with bookdown and R Markdown, and its source is fully available on GitHub. Machine Learning with Amazon SageMaker Cookbook Packt Publishing Ltd

InfoWorld is targeted to Senior IT professionals. Content is segmented into Channels and Topic Centers.

InfoWorld also celebrates people, companies, and projects.

InfoWorld Packt Publishing Ltd

Perform efficient fast text representation and classification with Facebook's fastText library Key Features Introduction to Facebook's fastText library for NLP Perform efficient word representations, sentence classification, vector representation Build better, more scalable solutions for text representation and classification Book Description Facebook's fastText library handles text representation and classification, used for Natural Language Processing (NLP). Most organizations have to deal with enormous amounts of text data on a daily basis, and gaining efficient data insights requires powerful NLP tools such as fastText. This book is your ideal introduction to fastText. You will learn how to create fastText models from the command line, without the need for complicated code. You will explore the algorithms that fastText is built on and how to use them for word representation and text classification. Next, you will use fastText in conjunction with other popular libraries and frameworks such as Keras, TensorFlow, and PyTorch. Finally, you will deploy fastText models to mobile devices. By the end of this book, you will have all the required knowledge to use fastText in your own applications at work or in projects. What you will learn Create models using the default command line options in fastText Understand the algorithms used in fastText to create word vectors Combine command line text transformation capabilities and the fastText library to implement a training, validation, and prediction pipeline Explore word representation and sentence classification using fastText Use Gensim and spaCy to load the vectors, transform, lemmatize, and perform other NLP tasks efficiently Develop a fastText NLP classifier using popular frameworks, such as Keras, Tensorflow, and PyTorch Who this book is for This book is for data analysts, data scientists, and machine learning developers who want to perform efficient word representation and sentence classification using Facebook's fastText library. Basic knowledge of Python programming is required.

Geospatial Data Analytics on AWS SAS Institute

Graph data closes the gap between the way humans and computers view the world. While computers rely on static rows and columns of data, people navigate and reason about life through relationships. This practical guide demonstrates how graph data brings these two approaches together. By working with concepts from graph theory, database schema, distributed systems, and data analysis, you ' ll arrive at a unique intersection known as graph thinking. Authors Denise Koessler Gosnell and Matthias Broecheler show data engineers, data scientists, and data analysts how to solve complex problems with graph databases. You ' ll explore templates for building with graph technology, along with examples that demonstrate how teams think about graph data within an application. Build an example application architecture with relational and graph technologies Use graph technology to build a Customer 360 application, the most popular graph data pattern today Dive into hierarchical data and troubleshoot a new paradigm that comes from working with graph data Find paths in graph data and learn why your trust in different paths motivates and informs your preferences Use collaborative filtering to design a Netflix-inspired recommendation system

Python Programming and Numerical Methods No Starch Press

Analyze and transform data efficiently with DuckDB, a versatile, modern, in-process SQL database Key Features Use DuckDB to rapidly load, transform, and query data across a range of sources and formats Gain practical experience using SQL, Python, and R to effectively analyze data Learn how open source tools and cloud services in the broader data ecosystem complement DuckDB ' s versatile capabilities Purchase of the print or Kindle book includes a free PDF eBook Book DescriptionDuckDB is a fast in-process analytical database. Its ease of use, versatile feature set, and powerful analytical capabilities make DuckDB a valuable addition to the data practitioner ' s toolkit. Getting Started with DuckDB offers a practical overview of DuckDB ' s fundamentals and guidance for effectively using its powerful capabilities. Through extensive hands-on examples, you ' ll learn how to use DuckDB to load, transform, and query a variety of data sources and formats, including CSV, JSON, and Parquet files, semi-structured data, remotely-hosted files, and external databases. You'll also find out how to leverage DuckDB's performance optimizations and friendly SQL enhancements. You'll explore how to use DuckDB ' s extensions for specialized applications, such as geospatial analysis and text search over document collections. In addition to working through examples in SQL, Python, and R, you ' ll also dive into using DuckDB for analyzing public datasets and discover the wider ecosystem of open-source tools and cloud services that supercharge DuckDB-powered workflows and applications. Whether you ' re a seasoned data practitioner or new to working with analytical data, this book will rapidly get you up to speed with DuckDB ' s versatile and powerful capabilities, enabling you to apply them in your analytical workflows and projects.What you will learn Understand the properties and applications of a columnar in-process database Use SQL to load, transform, and query a range of data formats Discover DuckDB's rich extensions and learn how to apply them Use nested data types to model semi-structured data and extract and model JSON data Integrate DuckDB into your Python and R analytical workflows Effectively leverage DuckDB's convenient SQL enhancements Explore the wider ecosystem and pathways for building DuckDB-powered data applications Who this book is for If you ' re interested in expanding your analytical toolkit, this book is for you. It will be particularly valuable for data analysts wanting to rapidly explore and query complex data, data and software engineers looking for a lean and versatile data processing tool, along with data scientists needing a scalable data manipulation library that integrates seamlessly with Python and R. You will get the most from this book if you have some familiarity with SQL and foundational database concepts, as well as exposure to a programming language such as Python or R.

Popular Photography CRC Press

Learn critical music production techniques used by professional producers to create a signature sound,

arranging, and mixing productions in FL Studio 21 from a Billboard charting JUNO nominated producer Purchase of the print or Kindle book includes a free PDF eBook Key Features Discover how top artists produce mind-blowing music and create ultimate records in FL Studio through this illustrated guide Build your skills in synchronizing songs placements and refining arrangements Master a professional workflow formula for producing chart-topping records Book DescriptionA digital audio workstation empowering both aspiring and seasoned producers to create original music compositions, FL Studio has not only advanced the culture of collaboration across several genres but has provided a creative outlet for up-and-coming artists worldwide. Achieving professional production prowess takes practice, market insight, and mentorship. This book explains how the author used FL Studio as a creative palette to build a successful career as a record producer, using specific techniques and workflow processes that only FL Studio can accommodate. You ' ll develop a Power User's mindset, create signature sounds using stock FL Studio One Shots, create top-level drum loops, learn about FL Studio's VST ' s, and approach arrangements from a practical and pop music perspective. This comprehensive guide covers everything from crafting and adding hypnotic melodies and chords, to mixing and mastering productions, and promoting those records to artists and companies, to take your career to the next level. By the end of this book, you ' ll be able to create original productions from scratch using FL Studio ' s virtual instruments and sound kits, mix and master the finished production, and arrange it using the Billboard-charting formula.What you will learn Acquaint yourself with FL Studio tools and maintain high-level workflow habits and frameworks Create ear-catching top-line melodies and advanced chord progressions Produce custom sound design and signature production sound Explore techniques to produce drum patterns and rhythm Uncover advanced music theory and piano roll production techniques Become a maestro in arranging productions for smash hits Learn unconventional mixing and mastering techniques Find tips on approaching, negotiating, and placing productions with artists, TV, advertising, and film companies Who this book is for If you are a music producer with intermediate knowledge of FL Studio or have experience using it, this guide will show you the author ' s personal tips and secrets for achieving Billboards success, landing major film and TV placements, and working with legendary artists. Tailored to experienced FL Studio users with a grasp on its fundamentals, this book assumes operational knowledge of FL Studio, including familiarity with its plugins and tools.

Popular Photography Packt Publishing Ltd

Learn dynamic programming with Julia to build apps for data analysis, visualization, machine learning, and the web Key FeaturesLeverage Julia's high speed and efficiency to build fast, efficient applicationsPerform supervised and unsupervised machine learning and time series analysisTackle problems concurrently and in a distributed environmentBook Description Julia offers the high productivity and ease of use of Python and R with the lightning-fast speed of C++. There ' s never been a better time to learn this language, thanks to its large-scale adoption across a wide range of domains, including fintech, biotech and artificial intelligence (AI). You will begin by learning how to set up a running Julia platform, before exploring its various built-in types. This Learning Path walks you through two important collection types: arrays and matrices. You ' ll be taken through how type conversions and promotions work, and in further chapters you'll study how Julia interacts with operating systems and other languages. You ' ll also learn about the use of macros, what makes Julia suitable for numerical and scientific computing, and how to run external programs. Once you have grasped the basics, this Learning Path goes on to how to analyze the Iris dataset using DataFrames. While building a web scraper and a web app, you ' ll explore the use of functions, methods, and multiple dispatches. In the final chapters, you'll delve into machine learning, where you'll build a book recommender system. By the end of this Learning Path, you ' ll be well versed with Julia and have the skills you need to leverage its high speed and efficiency for your applications. This Learning Path includes content from the following Packt products: Julia 1.0 Programming - Second Edition by Ivo BalbaertJulia Programming Projects by Adrian SalceanuWhat you will learnCreate your own types to extend the built-in type systemVisualize your data in Julia with plotting packagesExplore the use of built-in macros for testing and debuggingIntegrate Julia with other languages such as C, Python, and MATLABAnalyze and manipulate datasets using Julia and DataFramesDevelop and run a web app using Julia and the HTTP packageBuild a recommendation system using supervised machine learningWho this book is for If you are a statistician or data scientist who wants a quick course in the Julia programming language while building big data applications, this Learning Path is for you. Basic knowledge of mathematics and programming is a must.

A Power User's Guide to FL Studio 21 Packt Publishing Ltd

A hands-on guide to working on use cases helping you ingest, analyze, and serve insightful data from IoT as well as telemetry data sources using Azure Synapse Data Explorer Free PDF included with this book Key FeaturesAugment advanced analytics projects with your IoT and application dataExpand your existing Azure Synapse environments with unstructured dataBuild industry-level projects on integration, experimentation, and dashboarding with Azure SynapseBook Description Large volumes of data are generated daily from applications, websites, IoT devices, and other free-text, semi-structured data sources. Azure Synapse Data Explorer helps you collect, store, and analyze such data, and work with other analytical engines, such as Apache Spark, to develop advanced data science projects and maximize the value you extract from data. This book offers a comprehensive view of Azure Synapse Data Explorer, exploring not only the core scenarios of Data Explorer but also how it integrates within Azure Synapse. From data ingestion

to data visualization and advanced analytics, you'll learn to take an end-to-end approach to maximize the value of unstructured data and drive powerful insights using data science capabilities. With real-world usage scenarios, you'll discover how to identify key projects where Azure Synapse Data Explorer can help you achieve your business goals. Throughout the chapters, you'll also find out how to manage big data as part of a software as a service (SaaS) platform, as well as tune, secure, and serve data to end users. By the end of this book, you'll have mastered the big data life cycle and you'll be able to implement advanced analytical scenarios from raw telemetry and log data. What you will learnIntegrate Data Explorer pools with all other Azure Synapse servicesCreate Data Explorer pools with Azure Synapse Studio and Azure PortalIngest, analyze, and serve data to users using Azure Synapse pipelinesIntegrate Power BI and visualize data with Synapse StudioConfigure Azure Machine Learning integration in Azure SynapseManage cost and troubleshoot Data Explorer pools in Synapse AnalyticsSecure Synapse workspaces and grant access to Data Explorer poolsWho this book is for If you are a data engineer, data analyst, or business analyst working with unstructured data and looking to learn how to maximize the value of such data, this book is for you. If you already have experience working with Azure Synapse and want to incorporate unstructured data into your data science project, you'll also find plenty of useful information in this book. To maximize your learning experience, familiarity with data and performing simple queries using SQL or KQL is recommended. Basic knowledge of Python will help you get more from the examples.

Official Google Cloud Certified Professional Machine Learning Engineer Study Guide John Wiley & Sons

With this practical book, AI and machine learning practitioners will learn how to successfully build and deploy data science projects on Amazon Web Services. The Amazon AI and machine learning stack unifies data science, data engineering, and application development to help level upyour skills. This guide shows you how to build and run pipelines in the cloud, then integrate the results into applications in minutes instead of days. Throughout the book, authors Chris Fregly and Antje Barth demonstrate how to reduce cost and improve performance. Apply the Amazon AI and ML stack to real-world use cases for natural language processing, computer vision, fraud detection, conversational devices, and more Use automated machine learning to implement a specific subset of use cases with SageMaker Autopilot Dive deep into the complete model development lifecycle for a BERT-based NLP use case including data ingestion, analysis, model training, and deployment Tie everything together into a repeatable machine learning operations pipeline Explore real-time ML, anomaly detection, and streaming analytics on data streams with Amazon Kinesis and Managed Streaming for Apache Kafka Learn security best practices for data science projects and workflows including identity and access management, authentication, authorization, and more

Azure Data Scientist Associate Certification Guide O'Reilly Media

Python Programming and Numerical Methods: A Guide for Engineers and Scientists introduces programming tools and numerical methods to engineering and science students, with the goal of helping the students to develop good computational problem-solving techniques through the use of numerical methods and the Python programming language. Part One introduces fundamental programming concepts, using simple examples to put new concepts quickly into practice. Part Two covers the fundamentals of algorithms and numerical analysis at a level that allows students to quickly apply results in practical settings. Includes tips, warnings and "try this" features within each chapter to help the reader develop good programming practice Summaries at the end of each chapter allow for quick access to important information Includes code in Jupyter notebook format that can be directly run online

Natural Language Processing with Transformers, Revised Edition Packt Publishing Ltd

THE NEW YORK TIMES BESTSELLER Transform your life using the Bullet Journal Method, the revolutionary organisational system and worldwide phenomenon. The Bullet Journal Method will undoubtedly transform your life, in more ways than you can imagine! Hal Elrod, author of The Miracle Morning In his long-awaited first book, Ryder Carroll, the creator of the enormously popular Bullet Journal organisational system, explains how to use his method to: * TRACK YOUR PAST: using nothing more than a pen and paper, create a clear, comprehensive, and organised record of your thoughts and goals. * ORDER YOUR PRESENT: find daily calm by prioritising and minimising your workload and tackling your to-do list in a more mindful and productive way. * PLAN YOUR FUTURE: establish and appraise your short-term and long-term goals, plan more complex projects simply and effectively, and live your life with meaning and purpose. Like many of us, Ryder Carroll tried everything to get organised - countless apps, systems, planners, you name it. Nothing really worked. Then he invented his own simple system that required only pen and paper, which he found both effective and calming. He shared his method with a few friends, and before long he had a worldwide viral movement. The system combines elements of a wishlist, a to-do list, and a diary. It helps you identify what matters and set goals accordingly. By breaking long-term goals into small actionable steps, users map out an approachable path towards continual improvement, allowing them to stay focused despite the crush of incoming demands. But this is much more than a time management book. It's also a manifesto for what Ryder calls "intentional living": making sure that your beliefs and actions align. Even if you already use a Bullet Journal, this book gives you new exercises to become more calm and focused, new insights on how to prioritise well, and a new awareness of the power of analogue tools in a digital world. *** This book has been printed with three different colour designs, black, Nordic blue and emerald. We are unable to accept requests for a specific cover. The different covers will be assigned to orders at random. *** Beginning Database Design Solutions Packt Publishing Ltd Publisher ' s Note: Microsoft will stop supporting .NET 7 from May 2024. The newer 8th edition of the

book is available that covers .NET 8 (end-of-life November 2026) with C# 12 and EF Core. Purchase of the print or Kindle book includes a free PDF eBook Key Features Explore the newest additions to C# 11, the .NET 7 class libraries, and Entity Framework Core 7 Create professional websites and services with ASP.NET Core 7 and Blazor Build your confidence with step-by-step code examples and tips for best practices Book DescriptionExtensively revised to accommodate the latest features that come with C# 11 and .NET 7, this latest edition of our guide will get you coding in C# with confidence. You ' ll learn object-oriented programming, writing, testing, and debugging functions, implementing interfaces, and inheriting classes. Next, you ' ll take on .NET APIs for performing tasks like managing and querying data, working with the filesystem, and serialization. As you progress, you ' ll also explore examples of cross-platform projects you can build and deploy, such as websites and services using ASP.NET Core. Instead of distracting you with unnecessary graphical user interface code, the first eleven chapters will teach you about C# language constructs and many of the .NET libraries through simple console applications. Having mastered the basics, you ' ll then start building websites, web services, and browser apps. By the end of this book, you ' ll be able to create rich web experiences and have a solid grasp of object-oriented programming that you can build upon.What you will learn Build rich web experiences using Blazor, Razor Pages, the Model-View-Controller (MVC) pattern, and other features of ASP.NET Core Write, test, and debug functions Query and manipulate data using LINQ Integrate and update databases in your apps using Entity Framework Core models Build and consume powerful services using the latest technologies, including Web API and Minimal API For .NET 6 developers: C# 11 can be used with .NET 6, including features like raw string literals EF Core 7 targets .NET 6, so you can benefit from its new features like ExecuteUpdate and ExecuteDelete for more efficient data modifications Who this book is forThis book is primarily for beginners, but intermediate-level C# and .NET programmers who have worked with C# in the past and want to catch up with the changes made in the past few years will also find plenty of useful information in it. Prior exposure to C# or .NET is not a prerequisite, but you should have a general understanding of programming before you jump in. If you already have some C# and .NET skills and want to focus on developing apps, we recommend that you pick up Mark ' s other .NET book, Apps and Services with .NET 7, instead.

InfoWorld John Wiley & Sons

Discover all the essential design and architectural patterns in one place to help you rapidly build and deploy your modern data platform using AWS services Key Features Learn to build modern data platforms on AWS using data lakes and purpose-built data services Uncover methods of applying security and governance across your data platform built on AWS Find out how to operationalize and optimize your data platform on AWS Purchase of the print or Kindle book includes a free PDF eBook Book DescriptionMany IT leaders and professionals are adept at extracting data from a particular type of database and deriving value from it. However, designing and implementing an enterprise-wide holistic data platform with purpose-built data services, all seamlessly working in tandem with the least amount of manual intervention, still poses a challenge. This book will help you explore end-to-end solutions to common data, analytics, and AI/ML use cases by leveraging AWS services. The chapters systematically take you through all the building blocks of a modern data platform, including data lakes, data warehouses, data ingestion patterns, data consumption patterns, data governance, and AI/ML patterns. Using real-world use cases, each chapter highlights the features and functionalities of numerous AWS services to enable you to create a scalable, flexible, performant, and cost-effective modern data platform. By the end of this book, you ' ll be equipped with all the necessary architectural patterns and be able to apply this knowledge to efficiently build a modern data platform for your organization using AWS services.What you will learn Familiarize yourself with the building blocks of modern data architecture on AWS Discover how to create an end-to-end data platform on AWS Design data architectures for your own use cases using AWS services Ingest data from disparate sources into target data stores on AWS Build data pipelines, data sharing mechanisms, and data consumption patterns using AWS services Find out how to implement data governance using AWS services Who this book is for This book is for data architects, data engineers, and professionals creating data platforms. The book's use case — driven approach helps you conceptualize possible solutions to specific use cases, while also providing you with design patterns to build data platforms for any organization. It's beneficial for technical leaders and decision makers to understand their organization's data architecture and how each platform component serves business needs. A basic understanding of data & analytics architectures and systems is desirable along with beginner ' s level understanding of AWS Cloud.

Make Ink Fourth Estate

Expert, guidance for the Google Cloud Machine Learning certification exam In Google Cloud Certified Professional Machine Learning Study Guide, a team of accomplished artificial intelligence (AI) and machine learning (ML) specialists delivers an expert roadmap to AI and ML on the Google Cloud Platform based on new exam curriculum. With Sybex, you ' ll prepare faster and smarter for the Google Cloud Certified Professional Machine Learning Engineer exam and get ready to hit the ground running on your first day at your new job as an ML engineer. The book walks readers through the machine learning process from start to finish, starting with data, feature engineering, model training, and deployment on Google Cloud. It also discusses best practices on when to pick a custom model vs AutoML or pretrained models with Vertex AI platform. All technologies such as Tensorflow, Kubeflow, and Vertex AI are presented by way of real-world scenarios to help you apply the theory to practical examples and show you how IT professionals design, build, and operate secure ML cloud environments. The book also shows you how to: Frame ML problems and architect ML solutions from scratch Banish test anxiety by verifying and checking your progress with built-in self-assessments and other practical tools

Use the Sybex online practice environment, complete with practice questions and explanations, a glossary, objective maps, and flash cards A can't-miss resource for everyone preparing for the Google Cloud Certified Professional Machine Learning certification exam, or for a new career in ML powered by the Google Cloud Platform, this Sybex Study Guide has everything you need to take the next step in your career.

AWS Certified Machine Learning Study Guide Packt Publishing Ltd

InfoWorld is targeted to Senior IT professionals. Content is segmented into Channels and Topic Centers.

InfoWorld also celebrates people, companies, and projects.

[C# 11 and .NET 7 – Modern Cross-Platform Development Fundamentals](#) McFarland

Bring elasticity and innovation to Machine Learning and AI operations KEY FEATURES Coverage

includes a wide range of AWS AI and ML services to help you speedily get fully operational with ML.

Packed with real-world examples, practical guides, and expert data science methods for improving

AI/ML education on AWS. Includes ready-made, purpose-built models as AI services and proven

methods to adopt MLOps techniques. DESCRIPTION Using machine learning and artificial intelligence

(AI) in existing business processes has been successful. Even AWS's ML and AI services make it simple

and economical to conduct machine learning experiments. This book will show readers how to use the

complete set of AI and ML services available on AWS to streamline the management of their whole AI

operation and speed up their innovation. In this book, you'll learn how to build data lakes, build and

train machine learning models, automate MLOps, ensure maximum data reusability and reproducibility,

and much more. The applications presented in the book show how to make the most of several different

AWS offerings, including Amazon Comprehend, Amazon Rekognition, Amazon Lookout, and

AutoML. This book teaches you to manage massive data lakes, train artificial intelligence models, release

these applications into production, and track their progress in real-time. You will learn how to use the

pre-trained models for various tasks, including picture recognition, automated data extraction,

image/video detection, and anomaly detection. Every step of your Machine Learning and AI project's

development process is optimised throughout the book by utilising Amazon's pre-made, purpose-built

AI services. WHAT YOU WILL LEARN Learn how to build, deploy, and manage large-scale AI and

ML applications on AWS. Get your hands dirty with AWS AI services like SageMaker, Comprehend,

Rekognition, Lookout, and AutoML. Master data transformation, feature engineering, and model

training with Amazon SageMaker modules. Use neural networks, distributed learning, and deep

learning algorithms to improve ML models. Use AutoML, SageMaker Canvas, and Autopilot for

Model Deployment and Evaluation. Acquire expertise with Amazon SageMaker Studio, Jupyter

Server, and ML frameworks such as TensorFlow and MXNet. WHO THIS BOOK IS FOR Data

Engineers, Data Scientists, AWS and Cloud Professionals who are comfortable with machine learning

and the fundamentals of Python will find this book powerful. Familiarity with AWS would be helpful but

is not required. TABLE OF CONTENTS 1. Introducing the ML Workflow 2. Hydrating the Data Lake 3.

Predicting the Future With Features 4. Orchestrating the Data Continuum 5. Casting a Deeper Net

(Algorithms and Neural Networks) 6. Iteration Makes Intelligence (Model Training and Tuning) 7. Let

George Take Over (AutoML in Action) 8. Blue or Green (Model Deployment Strategies) 9. Wisdom at

Scale with Elastic Inference 10. Adding Intelligence with Sensory Cognition 11. AI for Industrial

Automation 12. Operationalized Model Assembly (MLOps and Best Practices)

Serverless ETL and Analytics with AWS Glue Packt Publishing Ltd

Quickly build and deploy machine learning models without managing infrastructure, and

improve productivity using Amazon SageMaker's capabilities such as Amazon SageMaker

Studio, Autopilot, Experiments, Debugger, and Model Monitor Key Features Build, train, and

deploy machine learning models quickly using Amazon SageMaker Analyze, detect, and receive

alerts relating to various business problems using machine learning algorithms and

techniques Improve productivity by training and fine-tuning machine learning models in

production Book Description Amazon SageMaker enables you to quickly build, train, and deploy

machine learning (ML) models at scale, without managing any infrastructure. It helps you focus

on the ML problem at hand and deploy high-quality models by removing the heavy lifting

typically involved in each step of the ML process. This book is a comprehensive guide for data

scientists and ML developers who want to learn the ins and outs of Amazon SageMaker. You'll

understand how to use various modules of SageMaker as a single toolset to solve the challenges

faced in ML. As you progress, you'll cover features such as AutoML, built-in algorithms and

frameworks, and the option for writing your own code and algorithms to build ML models. Later,

the book will show you how to integrate Amazon SageMaker with popular deep learning libraries

such as TensorFlow and PyTorch to increase the capabilities of existing models. You'll also

learn to get the models to production faster with minimum effort and at a lower cost. Finally,

you'll explore how to use Amazon SageMaker Debugger to analyze, detect, and highlight

problems to understand the current model state and improve model accuracy. By the end of this

Amazon book, you'll be able to use Amazon SageMaker on the full spectrum of ML workflows,

from experimentation, training, and monitoring to scaling, deployment, and automation. What

you will learn Create and automate end-to-end machine learning workflows on Amazon Web

Services (AWS) Become well-versed with data annotation and preparation techniques Use

AutoML features to build and train machine learning models with AutoPilot Create models using

built-in algorithms and frameworks and your own code Train computer vision and NLP models

using real-world examples Cover training techniques for scaling, model optimization, model

debugging, and cost optimization Automate deployment tasks in a variety of configurations using

SDK and several automation tools Who this book is for This book is for software engineers,

machine learning developers, data scientists, and AWS users who are new to using Amazon

SageMaker and want to build high-quality machine learning models without worrying about

infrastructure. Knowledge of AWS basics is required to grasp the concepts covered in this book

more effectively. Some understanding of machine learning concepts and the Python

programming language will also be beneficial.