

---

# Studio Notebook User Guide

When somebody should go to the book stores, search launch by shop, shelf by shelf, it is essentially problematic. This is why we present the book compilations in this website. It will very ease you to look guide Studio Notebook User Guide as you such as.

By searching the title, publisher, or authors of guide you truly want, you can discover them rapidly. In the house, workplace, or perhaps in your method can be every best area within net connections. If you object to download and install the Studio Notebook User Guide, it is agreed easy then, in the past currently we extend the colleague to buy and make bargains to download and install Studio Notebook User Guide in view of that simple!



fastText Quick Start  
Guide Packt Publishing  
Ltd  
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 up your 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

Modern Data Architecture on AWS IBM Redbooks

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. **Applied Machine Learning for Healthcare and Life Sciences Using AWS** John Wiley & Sons 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 Features Explore the various AWS services that make automated machine learning easier Recognize the role of DevOps and MLOps methodologies in pipeline automation Get acquainted with additional AWS services such as Step Functions, MWAA, and more to overcome automation challenges Book 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 learn Employ SageMaker Autopilot and Amazon SageMaker SDK to automate the

---

machine learning process Understand how to use AutoGluon to automate complicated model building tasks Use the AWS CDK to codify the machine learning process Create, deploy, and rebuild a CI/CD pipeline on AWS Build an ML workflow using AWS Step Functions and the Data Science SDK Leverage 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 process Who 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.

### **Serverless ETL and Analytics with AWS Glue**

BPB Publications

Explore Azure AI Platform KEY FEATURES ? Easy-to-follow tutorial for getting started with the Azure AI platform. ? Integrated platform for developing, deploying, and managing AI apps. ? Includes real-world scenarios and use-cases to fully explore Azure AI

Platform. DESCRIPTION Microsoft Azure AI A Beginner's Guide explains the fundamentals of Azure AI and some more advanced topics. The sole objective of the book is to provide hands-on experience working with the various services, APIs, and tools available in the Azure AI Platform. This book begins by discussing the fundamentals of the Azure AI platform and the essential principles behind the Azure AI ecosystem and services. Readers will become familiar with the essential services, use cases, and examples provided by Azure AI Platform and Services, including Azure Cognitive

---

Services, Azure Computer Vision, Azure Applied AI Services, and Azure Machine Learning. The author focuses on teaching how to utilize Azure Cognitive services to construct intelligent apps, including Image Processing, Object Detection, Text Recognition, OCR, Spatial Analysis, and Face Recognition using Computer Vision. Readers can investigate Azure Applied AI Services, including Form Recognizer, Metrics Advisor, Cognitive Search, Immersive Reader, Video Analyzer, and Azure Bot Service. Bot Framework and the Bot Framework Emulator will be

explored in further detail, and how they can be used in AI applications to improve their conversational user interfaces. With Azure Machine Learning Studio, you will also learn to incorporate machine learning into your enterprise-level applications. WHAT YOU WILL LEARN ? Get familiar with Azure AI Platform and the cognitive capabilities of Azure. ? Learn to create apps that can process photos, detect faces, and detect objects. ? Utilize OCR, handwriting recognition, and spatial analysis in your development. ? Learn about Azure AI services like Form Recognizer, Metrics Advisor, Cognitive Search, Azure

Immersive Reader, and Video Analyzer. ? Try out several NLP applications with the Azure BOT framework. WHO THIS BOOK IS FOR This book teaches AI developers, machine learning engineers, .NET developers, and architects how to swiftly develop intelligent applications utilizing the Azure AI Platform. Knowledge of .NET or .NET Core is strongly advised to get the most out of the book. TABLE OF CONTENTS 1 .Azure AI Platform and Services 2. Azure Computer Vision - Image Analysis, Processing, Content Moderation, Object and Face Detection 3. Computer Vision -

---

Text Recognition, Optical Character Recognition, Spatial Analysis 4. Azure Cognitive Services - Custom Applications leveraging Decision, Language, Speech, Web Search 5. Azure Applied AI Services 6. Azure Applied AI Services -BOTS- A Brief Introduction 7. Machine Learning-Infusing ML in Custom Applications using ML.NET 8. Machine Learning - Using Azure ML Studio *Reference Guide to Russian Literature* BPB Publications It's all in the name: Learn You a Haskell for Great

Good! is a hilarious, illustrated guide to this complex functional language. Packed with the author's original artwork, pop culture references, and most importantly, useful example code, this book teaches functional fundamentals in a way you never thought possible. You'll start with

the kid stuff: basic syntax, recursion, types and type classes. Then once you've got the basics down, the real black belt master-class begins: you'll learn to use applicative functors, monads, zippers, and all the other mythical Haskell constructs you've only read about in storybooks. As you



---

work your way through the author's imaginative (and occasionally insane) examples, you'll learn to:

- Laugh in the face of side effects as you wield purely functional programming techniques
- Use the magic of Haskell's "laziness" to play with infinite sets of data
- Organize your programs by

creating your own types, type classes, and modules -Use Haskell's elegant input/output system to share the genius of your programs with the outside world Short of eating the author's brain, you will not find a better way to learn this powerful language than reading Learn You a Haskell for Great Good!

## **The Definitive Guide to Google Vertex AI**

McFarland

Build efficient data lakes that can scale to virtually unlimited size using AWS Glue Key Features Book Description Organizations these days have gravitated toward services such as AWS Glue that undertake undifferentiated heavy lifting and provide serverless Spark, enabling you to create and manage data lakes in a serverless fashion. This guide shows you how AWS Glue

---

can be used to solve real-world problems along with helping you learn about data processing, data integration, and building data lakes. Beginning with AWS Glue basics, this book teaches you how to perform various aspects of data analysis such as ad hoc queries, data visualization, and real-time analysis using this service. It also provides a walk-through of CI/CD for AWS Glue and how to shift left on quality using automated regression tests. You'll find out how data security aspects such as access control, encryption, auditing, and networking are implemented, as well as getting to grips with useful techniques such as picking the right file format, compression, partitioning, and bucketing. As you advance, you'll discover AWS Glue features such as crawlers, Lake Formation, governed tables, lineage, DataBrew, Glue Studio, and custom connectors. The concluding chapters help you to understand various performance tuning, troubleshooting, and monitoring options. By the end of this AWS book, you'll be able to create, manage, troubleshoot, and deploy ETL pipelines using AWS Glue. What you will learn Apply various AWS Glue features to manage and create data lakes Use Glue DataBrew and Glue Studio for data preparation Optimize data layout in cloud

---

storage to accelerate analytics workloads  
Manage metadata including database, table, and schema definitions  
Secure your data during access control, encryption, auditing, and networking  
Monitor AWS Glue jobs to detect delays and loss of data  
Integrate Spark ML and SageMaker with AWS Glue to create machine learning models  
Who this book is for  
ETL developers, data engineers, and data analysts

## **Beginning Database**

**Design Solutions** No Starch Press  
InfoWorld is targeted to Senior IT professionals. Content is segmented into Channels and Topic Centers. InfoWorld also celebrates people, companies, and projects.

**Geospatial Data Analytics on AWS** John Wiley & Sons  
"The pigments he concocts from these humble beginnings are as fun to make as they are eye-opening to work with . . .

the world never quite looks the same."  
—MarthaStewart.com  
A 2018 Best Book of the Year—The Guardian  
The Toronto Ink Company was founded in 2014 by designer and artist Jason Logan as a citizen science experiment to make eco-friendly, urban ink from street-harvested pigments.  
In *Make Ink*, Logan delves into the history of inkmaking and the science of distilling pigment

---

from the natural world. Readers will learn how to forage for materials such as soot, rust, cigarette butts, peach pits, and black walnut, then how to mix, test, and transform these ingredients into rich, vibrant inks that are sensitive to both place and environment. Organized by color, and featuring lovely minimalist photography

throughout, *Make Ink* combines science, art, and craft to instill the basics of ink making and demonstrate the beauty and necessity of engaging with one of mankind's oldest tools of communication. "Logan demystifies the process, encouraging experimentation and taking a fresh look at urban environments." —NPR "The book is full of inspiration and takes

a lot of the mystery out of ink making, at least at its simplest level. And it also reminds me why I love ink—any ink or liquid color as much as I do." —The Well-Appointed Desk "Quite a few recipes . . . that use color from the kitchen: carrots, black beans, blueberries, turmeric, and onion skins all make beautiful ink colors." —Design Observer "Make Ink

---

opens up about supervised and  
methods, providing an unsupervised  
open source guide to machine learning  
DIY ink." –CityLab and time series  
*Make Ink* Abrams analysisTackle  
Learn dynamic problems  
programming with concurrently and in  
Julia to build apps a distributed  
for data analysis, environmentBook  
visualization, Description Julia  
machine learning, offers the high  
and the web Key productivity and  
FeaturesLeverage ease of use of  
Julia's high speed Python and R with  
and efficiency to the lightning-fast  
build fast, speed of C++.  
efficient There's never been  
applicationsPerform 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 scientific  
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

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:	and	or data scientist
Julia 1.0	debuggingIntegrate	who wants a quick
Programming -	Julia with other	course in the Julia
Second Edition by	languages such as	programming
Ivo BalbaertJulia	C, Python, and	language while
Programming	MATLABAnalyze and	building big data
Projects by Adrian	manipulate datasets	applications, this
SalceanuWhat you	using Julia and	Learning Path is
will learnCreate	DataFramesDevelop	for you. Basic
your own types to	and run a web app	knowledge of
extend the built-in	using Julia and the	mathematics and
type	HTTP packageBuild a	programming is a
systemVisualize	recommendation	must.
your data in Julia	system using	<u><a href="#">AWS Certified Machine</a></u>
with plotting	supervised machine	<u><a href="#">Learning Study Guide</a></u>
packagesExplore the	learningWho this	John Wiley & Sons
use of built-in	book is for If you	R Markdown: The
macros for testing	are a statistician	Definitive Guide is
		the first official

---

book authored by the core R Markdown developers that provides a comprehensive and accurate reference to the R Markdown ecosystem. With R Markdown, you can easily create reproducible data analysis reports, presentations, dashboards, interactive applications, books, dissertations, websites, and journal articles, while enjoying the simplicity of Markdown and the great power of R and other languages. In this book, you will learn Basics: Syntax of Markdown and R code chunks, how to generate figures and tables, and how to use other computing languages Built-in output formats of R Markdown: PDF/HTML/Word/RTF/Markdown documents and ioslides/Slidy/Beamer/PowerPoint presentations Extensions and applications: Dashboards, Tufte handouts, xaringan/reveal.js presentations, websites, books, journal articles, and interactive tutorials Advanced topics: Parameterized reports, HTML widgets, document templates, custom output formats, and Shiny documents. Yihui Xie is a software engineer at RStudio. He has authored and co-authored several R packages, including knitr, rmarkdown, bookdown, blogdown, shiny, xaringan, and animation. He has published three other books, Dynamic Documents with R and knitr, bookdown:



---

Authoring Books and Technical Documents with R Markdown, and blogdown: Creating Websites with R Markdown. J.J. Allaire is the founder of RStudio and the creator of the RStudio IDE. He is an author of several packages in the R Markdown ecosystem including rmarkdown, flexdashboard, learnr, and radix. Garrett Golemund is the co-author of R for Data Science and author of Hands-On Programming with R. He wrote the lubridate R package and

works for RStudio as an advocate who trains engineers to do data science with R and the Tidyverse.

[Learn You a Haskell for Great Good!](#)

John Wiley & Sons 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 for If 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. **Pragmatic AI** CRC Press

Succeed on the AWS Machine Learning exam or in your next job as a machine learning specialist on the AWS Cloud platform with this hands-on guide As the most popular cloud service in the world today, Amazon Web Services offers a wide range of opportunities for those interested in the development and deployment of artificial intelligence and machine learning business solutions. The AWS Certified Machine Learning Study Guide: Specialty (MLS-C01)

---

Exam delivers hyper-focused, authoritative instruction for anyone considering the pursuit of the prestigious Amazon Web Services Machine Learning certification or a new career as a machine learning specialist working within the AWS architecture. From exam to interview to your first day on the job, this study guide provides the domain-by-domain specific knowledge you need to build, train, tune, and deploy machine learning models with the AWS Cloud. And with the practice exams and assessments, electronic flashcards, and supplementary online resources that accompany this Study Guide, you'll be prepared for success in every subject area covered by the exam. You'll also find: An intuitive and organized layout perfect for anyone taking the exam for the first time or seasoned professionals seeking a refresher on machine learning on the AWS Cloud. And with the recognized certification that unlocks countless career opportunities in machine learning and data science. Access to the Sybex online learning resources and test bank, with chapter review questions, a full-length practice exam, hundreds of electronic flashcards, and a glossary of key terms. AWS Certified Machine Learning Specialty (MLS-C01) Exam is an indispensable guide for anyone seeking to prepare themselves for

---

success on the AWS Certified Machine Learning Specialty exam or for a job interview in the field of machine learning, or who wishes to improve their skills in the field as they pursue a career in AWS machine learning.

Getting Started with DuckDB Addison-Wesley Professional

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. Description DuckDB 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. Python Programming and Numerical Methods Packt Publishing Ltd Track-by-track, album-by-album guide to the ultimate cult act The Fall's 25-year recording career. **Official Gazette of the United States**

**Patent and Trademark Office** Helter Skelter Publishing  
A concise introduction to database design concepts, methods, and techniques in and out of the cloud In the newly revised second edition of *Beginning Database Design Solutions: Understanding and Implementing Database Design Concepts for the Cloud and Beyond, Second Edition*, award-winning programming



---

instructor and mathematician Rod Stephens delivers an easy-to-understand guide to designing and implementing databases both in and out of the cloud. Without assuming any prior database design knowledge, the author walks you through the steps you'll need to take to understand, analyze, design, and build databases. In the book, you'll find clear coverage of foundational database concepts along with hands-on examples that help you practice important techniques so you can apply them to your own database designs, as well as:

Downloadable source code that illustrates the concepts discussed in the book

Best practices for reliable, platform-agnostic database design Strategies for digital transformation driven by universally accessible database design An essential resource for database administrators, data management specialists, and database developers seeking expertise in relational, NoSQL, and hybrid database design both in and out of the cloud,

Beginning Database Design Solutions is a hands-on guide ideal for students and practicing professionals alike.

Natural Language

---

Processing with Transformers, Revised Edition Packt Publishing Ltd  
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 Description  
Many 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.

### **A Reference Guide to Television's**

**Bonanza** 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 single toolset to learning libraries high-quality models solve the such as TensorFlow by removing the challenges faced in and PyTorch to heavy lifting ML. As you increase the typically involved progress, you'll capabilities of in each step of the cover features such existing models. ML process. This as AutoML, built-in You'll also learn book is a algorithms and to get the models comprehensive guide frameworks, and the to production for data scientists option for writing faster with minimum and ML developers your own code and effort and at a who want to learn algorithms to build lower cost. the ins and outs of ML models. Later, Finally, you'll Amazon SageMaker. the book will show explore how to use You'll understand you how to Amazon SageMaker how to use various integrate Amazon Debugger to modules of SageMaker with analyze, detect, SageMaker as a popular deep 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

learnCreate and automate end-to-end machine learning workflows on Amazon Web Services (AWS)Become well-versed with data annotation and preparation techniquesUse AutoML features to build and train machine learning models with AutoPilotCreate models using built-in algorithms and frameworks and your

own codeTrain computer vision and NLP models using real-world examplesCover training techniques for scaling, model optimization, model debugging, and cost optimizationAutomate deployment tasks in a variety of configurations using SDK and several automation toolsWho 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. *R Data Mining* 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.

Learning SAS by Example Packt Publishing Ltd Implement machine learning pipelines with Google Cloud Vertex AI Key Features Understand the role of an AI platform and MLOps practices in machine learning projects Get acquainted with Google Vertex AI tools and offerings that help accelerate the

---

creation of end-to-end ML solutions. Implement Vision, NLP, and recommendation-based real-world ML models on Google Cloud Platform. Purchase of the print or Kindle book includes a free PDF eBook. Book Description While AI has become an integral part of every organization today, the development of

large-scale ML solutions and management of complex ML workflows in production continue to pose challenges for many. Google's unified data and AI platform, Vertex AI, directly addresses these challenges with its array of MLOps tools designed for overall workflow management. This book is a

comprehensive guide that lets you explore Google Vertex AI's easy-to-advanced level features for end-to-end ML solution development. Throughout this book, you'll discover how Vertex AI empowers you by providing essential tools for critical tasks, including data management, model building, large-scale



---

experimentations, leveraging key metadata logging, technologies such as Vision, NLP, model deployments, as generative AI, and recommendation systems. You'll learn how to harness the full potential of Vertex AI for developing and deploying no-code, low-code, or fully customized ML solutions. This book takes a hands-on approach to developing and deploying some real-world ML solutions on Google Cloud, leveraging key technologies such as Vision, NLP, generative AI, and recommendation systems. Additionally, this book covers pre-built and turnkey solution offerings as well as guidance on seamlessly integrating them into your ML workflows. By the end of this book, you'll have the confidence to develop and deploy large-scale production-grade ML solutions using the MLOps tooling and best practices from Google. What you will learn

- Understand the ML lifecycle, challenges, and importance of MLOps
- Get started with ML model development quickly using Google Vertex AI
- Manage datasets, artifacts, and

---

experiments Develop  
no-code, low-code,  
and custom AI  
solution on Google  
Cloud Implement  
advanced model  
optimization  
techniques and  
tooling Understand  
pre-built and  
turnkey AI solution  
offerings from  
Google Build and  
deploy custom ML  
models for real-  
world applications  
Explore the latest  
generative AI tools

within Vertex AI Who  
this book is for If  
you are a machine  
learning  
practitioner who  
wants to learn end-  
to-end ML solution  
development on  
Google Cloud  
Platform using  
MLOps best  
practices and tools  
offered by Google  
Vertex AI, this is  
the book for you.

**SQL Query Design  
Patterns and Best  
Practices** Packt

Publishing Ltd  
This IBM® Redpaper  
publication describes  
how to deploy Red Hat  
OpenShift V4.3 on IBM  
Power Systems servers.  
This book presents  
reference  
architectures for  
deployment, initial  
sizing guidelines for  
server, storage, and  
IBM Cloud® Paks.  
Moreover, this  
publication delivers  
information about  
initial supported  
Power System  
configurations for Red  
Hat OpenShift V4.3  
deployment (bare

---

metal, IBM PowerVM® LE Systems. This book also LPARs, and others). provides technical content to transfer how-to skills to the Red Hat OpenShift V4.3 support teams, and provide start solution guidance to guidelines and the sales team. This recommended practices book compliments the for implementing it on documentation that is Power Systems and available at IBM completing it with the Knowledge Center, and supported IBM Cloud also aligns with the Paks. The publication educational offerings addresses topics for that are provided by developers, IT the IBM Systems architects, IT Technical Education specialists, sellers, (SSE). and anyone who wants to implement a Red Hat OpenShift V4.3 and IBM Cloud Paks on IBM Power