



Computer Vision on AWS : Build and Deploy Real-World CV Solutions with Amazon Rekognition, Lookout for Vision, and SageMaker

Mullennex, Lauren

Monografía

Develop scalable computer vision solutions for real-world business problems and discover scaling, cost reduction, security, and bias mitigation best practices with AWS AI/ML services Purchase of the print or Kindle book includes a free PDF eBook Key Features Learn how to quickly deploy and automate end-to-end CV pipelines on AWS Implement design principles to mitigate bias and scale production of CV workloads Work with code examples to master CV concepts using AWS AI/ML services Book Description Computer vision (CV) is a field of artificial intelligence that helps transform visual data into actionable insights to solve a wide range of business challenges. This book provides prescriptive guidance to anyone looking to learn how to approach CV problems for quickly building and deploying production-ready models. You'll begin by exploring the applications of CV and the features of Amazon Rekognition and Amazon Lookout for Vision. The book will then walk you through real-world use cases such as identity verification, real-time video analysis, content moderation, and detecting manufacturing defects that'll enable you to understand how to implement AWS AI /ML services. As you make progress, you'll also use Amazon SageMaker for data annotation, training, and deploying CV models. In the concluding chapters, you'll work with practical code examples, and discover best practices and design principles for scaling, reducing cost, improving the security posture, and mitigating bias of CV workloads. By the end of this AWS book, you'll be able to accelerate your business outcomes by building and implementing CV into your production environments with the help of AWS AI/ML services. What you will learn Apply CV across industries, including e-commerce, logistics, and media Build custom image classifiers with Amazon Rekognition Custom Labels Create automated end-to-end CV workflows on AWS Detect product defects on edge devices using Amazon Lookout for Vision Build, deploy, and monitor CV models using Amazon SageMaker Discover best practices for designing and evaluating CV workloads Develop an AI governance strategy across the entire machine learning life cycle Who this book is for If you are a machine learning engineer or data scientist looking to discover best practices and learn how to build comprehensive CV solutions on AWS, this book is for you. Knowledge of AWS basics is required to grasp the concepts covered in this book more effectively. A solid understanding of machine learning concepts and the Python programming language will also be beneficial

Título: Computer Vision on AWS Build and Deploy Real-World CV Solutions with Amazon Rekognition, Lookout for Vision, and SageMaker

Edición: 1st ed

Editorial: Birmingham Packt Publishing, Limited 2023 2023

Descripción física: 1 online resource (324 pages)

Contenido: Cover -- Title Page -- Copyright and Credits -- Contributors -- Table of Contents -- Preface -- Part 1: Introduction to CV on AWS and Amazon Rekognition -- Chapter 1: Computer Vision Applications and AWS AI /ML Services Overview -- Technical requirements -- Understanding CV -- CV architecture and applications -- Data processing and feature engineering -- Data labeling -- Solving business challenges with CV -- Contactless check-in and checkout -- Video analysis -- Content moderation -- CV at the edge -- Exploring AWS AI/ML services -- AWS AI services -- Amazon SageMaker -- Setting up your AWS environment -- Creating an Amazon SageMaker Jupyter notebook instance -- Summary -- Chapter 2: Interacting with Amazon Rekognition -- Technical requirements -- The Amazon Rekognition console -- Using the Label detection demo -- Examining the API request -- Examining the API response -- Other demos -- Monitoring Amazon Rekognition -- Quick recap -- Detecting Labels using the API -- Uploading the images to S3 -- Initializing the boto3 client -- Detect the Labels -- Using the Label information -- Using bounding boxes -- Quick recap -- Cleanup -- Summary -- Chapter 3: Creating Custom Models with Amazon Rekognition Custom Labels -- Technical requirements -- Introducing Amazon Rekognition Custom Labels -- Benefits of Amazon Rekognition Custom Labels -- Creating a model using Rekognition Custom Labels -- Deciding the model type based on your business goal -- Creating a model -- Improving the model -- Starting your model -- Analyzing an image -- Stopping your model -- Building a model to identify Packt's logo -- Step 1 - Collecting your images -- Step 2 - Creating a project -- Step 3 - Creating training and test datasets -- Step 4 - Adding labels to the project -- Step 5 - Drawing bounding boxes on your training and test datasets -- Step 6 - Training your model Validating that the model works -- Step 1 - Starting your model -- Step 2 - Analyzing an image with your model -- Step 3 - Stopping your model -- Summary -- Part 2: Applying CV to Real-World Use Cases -- Chapter 4: Using Identity Verification to Build a Contactless Hotel Check-In System -- Technical requirements -- Prerequisites -- Creating the image bucket -- Uploading the sample images -- Creating the profile table -- Introducing collections -- Creating a collection -- Describing a collection -- Deleting a collection -- Quick recap -- Describing the user journeys -- Registering a new user -- Authenticating a user -- Registering a new user with an ID card -- Updating the user profile -- Implementing the solution -- Checking image quality -- Indexing face information -- Search existing faces -- Quick recap -- Supporting ID cards -- Reading an ID card -- Using the CompareFaces API -- Quick recap -- Guidance for identity verification on AWS -- Solution overview -- Deployment process -- Cleanup -- Summary -- Chapter 5: Automating a Video Analysis Pipeline -- Technical requirements -- Creating the video bucket -- Uploading content to Amazon S3 -- Creating the person-tracking topic -- Subscribing a message queue to the person-tracking topic -- Creating the person-tracking publishing role -- Setting up IP cameras -- Quick recap -- Using IP cameras -- Installing OpenCV -- Installing additional modules -- Connecting with OpenCV -- Viewing the frame -- Uploading the frame -- Reporting frame metrics -- Quick recap -- Using the PersonTracking API -- Uploading the video to Amazon S3 -- Using the StartPersonTracking API -- Receiving the completion notification -- Using the GetPersonTracking API -- Reviewing the GetPersonTracking response -- Viewing the frame -- Quick recap -- Summary -- Chapter 6: Moderating Content with AWS AI Services -- Technical requirements Moderating images -- Using the DetectModerationLabels API -- Using top-level categories -- Using secondary-level categories -- Putting it together -- Quick recap -- Moderating videos -- Creating the supporting resources -- Finding the resource ARNs -- Uploading the sample video to Amazon S3 -- Using the StartContentModeration API -- Examining the completion notification -- Using the GetContentModeration API -- Quick recap -- Using AWS Lambda to automate the workflow -- Implement the Start Analysis Handler -- Implementing the Get Results Handler -- Publishing function changes -- Experiment with the end-to-end -- Summary -- Part 3: CV at the edge -- Chapter 7: Introducing Amazon Lookout for Vision -- Technical requirements -- Introducing Amazon Lookout for Vision -- The benefits of Amazon Lookout for Vision -- Creating a model using Amazon Lookout for Vision -- Choosing the model type based on your business goals -- Creating a model -- Starting your model -- Analyzing an image -- Stopping your model -- Building a model to identify damaged pills -- Step 1 - collecting your images -- Step 2 - creating a project -- Step 3 - creating the training and test datasets -- Step 4 - verifying the dataset -- Step 5 - training your model -- Validating it works -- Step 1 - trial detection -- Step 2 -

starting your model -- Step 3 - analyzing an image with your model -- Step 4 - stopping your model -- Summary -- Chapter 8: Detecting Manufacturing Defects Using CV at the Edge -- Technical requirements -- Understanding ML at the edge -- Deploying a model at the edge using Lookout for Vision and AWS IoT Greengrass -- Step 1 - Launch an Amazon EC2 instance -- Step 2 - Create an IAM role and attach it to an EC2 instance -- Step 3 - Install AWS IoT Greengrass V2 -- Step 4 - Upload training and test datasets to S3 -- Step 5 - Create a project Step 6 - Create training and test datasets -- Step 7 - Train the model -- Step 8 - Package the model -- Step 9 - Configure IoT Greengrass IAM permissions -- Step 10 - Deploy the model -- Step 11 - Run inference on the model -- Step 12 - Clean up resources -- Summary -- Part 4: Building CV Solutions with Amazon SageMaker -- Chapter 9: Labeling Data with Amazon SageMaker Ground Truth -- Technical requirements -- Introducing Amazon SageMaker Ground Truth -- Benefits of Amazon SageMaker Ground Truth -- Automated data labeling -- Labeling Packt logos in images using Amazon SageMaker Ground Truth -- Step 1 - collect your images -- Step 2 - create a labeling job -- Step 3 - specify the job details -- Step 4 - specify worker details -- Step 5 - providing labeling instructions -- Step 6 - start labeling -- Step 7 - output data -- Importing the labeled data with Rekognition Custom Labels -- Step 1 - create the project -- Step 2 - create training and test datasets -- Step 3 - model training -- Summary -- Chapter 10: Using Amazon SageMaker for Computer Vision -- Technical requirements -- Fetching the LabelMe-12 dataset -- Installing TensorFlow 2.0 -- Installing matplotlib -- Using the built-in image classifier -- Upload the dataset to Amazon S3 -- Prepare the job channels -- Start the training job -- Monitoring and troubleshooting -- Quick recap -- Handling binary metadata files -- Declaring the Label class -- Reading the annotations file -- Declaring the Annotation class -- Validate parsing the file -- Restructure the files -- Load the dataset -- Quick recap -- Summary -- Part 5: Best Practices for Production-Ready CV Workloads -- Chapter 11: Integrating Human-in-the-Loop with Amazon Augmented AI (A2I) -- Technical requirements -- Introducing Amazon A2I -- Core concepts of Amazon A2I -- Learning how to build a human review workflow Creating a labeling workforce -- Setting up an A2I human review workflow or flow definition -- Initiating a human loop -- Leveraging Amazon A2I with Amazon Rekognition to review images -- Step 1 - Collecting your images -- Step 2 - Creating a work team -- Step 3 - Creating a human review workflow -- Step 4 - Starting a human loop -- Step 5 - Checking the human loop status -- Step 6 - Reviewing the output data -- Summary -- Chapter 12: Best Practices for Designing an End-to-End CV Pipeline -- Defining a problem that CV can solve and processing data -- Developing a CV model -- Training -- Evaluating -- Tuning -- Deploying and monitoring a CV model -- Shadow testing -- A/B testing -- Blue/Green deployment strategy -- Monitoring -- Developing an MLOps strategy -- SageMaker MLOps features -- Workflow automation tools -- Using the AWS Well-Architected Framework -- Cost optimization -- Operational excellence -- Reliability -- Performance efficiency -- Security -- Sustainability -- Summary -- Chapter 13: Applying AI Governance in CV -- Understanding AI governance -- Defining risks, documentation, and compliance -- Data risks and detecting bias -- Auditing, traceability, and versioning -- Monitoring and visibility -- MLOps -- Responsibilities of business stakeholders -- Applying AI governance in CV -- Types of biases -- Mitigating bias in identity verification workflows -- Using Amazon SageMaker for governance -- ML governance capabilities with Amazon SageMaker -- Amazon SageMaker Clarify for explainable AI -- Summary -- Index -- Other Books You May Enjoy

ISBN: 1-80324-820-3

Materia Título preferido: Amazon Web Services

Materia: Computer vision- Computer programs Informática en la nube Servicios Web Artificial intelligence

Autores: Bachmeier, Nate Rao, Jay

Enlace a formato físico adicional: Print version Mullenex, Lauren. Computer Vision on AWS Birmingham : Packt Publishing, Limited,c2023

Baratz Innovación Documental

- Gran Vía, 59 28013 Madrid
- (+34) 91 456 03 60
- informa@baratz.es