

Webinar Series on Remote Learning

TITLE

Detect, Analyze, and Compare Faces with Amazon Rekognition using AWS Educate Classrooms

BRIEF SYNOPSIS

Amazon Rekognition is a deep learning-based image and video analysis service. In this tutorial, you will learn about Amazon Rekognition and how to use the face recognition features in Amazon Rekognition by using the AWS Management Console.

In the webinar we cover:

- A brief introduction to Amazon Rekognition and some use cases
- A live demo of Amazon Rekognition
- Learn how to analyze an image using the AWS Management Console
- Compare images and analyze if they are the same

SPEAKER

Ryan Little

SPEAKER BIO:

Technical Program Manager, AWS WWPS

DURATION

29 min 37 sec

DURATION

29 min 37 sec

TIMESTAMPS

Detect, Analyze, and Compare Faces with Amazon Rekognition using AWS Educate Classrooms - Part 1

[0:00](#) – Speaker Introductions

[7:49](#) – The reach of Machine Learning

[16:08](#) – The AWS ML Stack

[17:45](#) – Introduction to Amazon Rekognition

[20:25](#) – Amazon Rekognition Use Cases

[26:17](#) – Amazon Rekognition Tutorial

TIMESTAMPS

Detect, Analyze, and Compare Faces with Amazon Rekognition using AWS Educate Classrooms - Part 2

[0:00](#) – How to analyze an image using the AWS Management Console

[6:05](#) – Compare images and analyze if they are the same

[11:27](#) – AWS Educate for Educators

Webinar Series on Remote Learning

If you enjoyed this content, be sure to check out our live and on-demand webinars [here](#).

FAQs

Q: What is Amazon Rekognition?

A: Amazon Rekognition is a service that makes it easy to add powerful visual analysis to your applications. Rekognition Image lets you easily build powerful applications to search, verify, and organize millions of images. Rekognition Video lets you extract motion-based context from stored or live stream videos and helps you analyze them.

Rekognition Image is an image recognition service that detects objects, scenes, and faces; extracts text; recognizes celebrities; and identifies inappropriate content in images. It also allows you to search and compare faces. Rekognition Image is based on the same proven, highly scalable, deep learning technology developed by Amazon's computer vision scientists to analyze billions of images daily for Prime Photos.

Rekognition Image uses deep neural network models to detect and label thousands of objects and scenes in your images, and we are continually adding new labels and facial recognition features to the service. With Rekognition Image, you only pay for the images you analyze and the face metadata you store.

Rekognition Video is a video recognition service that detects activities; understands the movement of people in frame; and recognizes objects, celebrities, and inappropriate content in videos stored in Amazon S3 and live video streams from Acuity. Rekognition Video detects persons and tracks them through the video even when their faces are not visible, or as the whole person might go in and out of the scene. For example, this could be used in an application that sends a real-time notification when someone delivers a package to your door. Rekognition Video allows you also to index metadata like objects, activities, scene, celebrities, and faces that make video search easy.

FAQs

Q: What is deep learning?

A: Deep learning is a sub-field of Machine Learning and a significant branch of Artificial Intelligence. It aims to infer high-level abstractions from raw data by using a deep graph with multiple processing layers composed of multiple linear and non-linear transformations. Deep learning is loosely based on models of information processing and communication in the brain. Deep learning replaces handcrafted features with ones learned from very large amounts of annotated data. Learning occurs by iteratively estimating hundreds of thousands of parameters in the deep graph with efficient algorithms. Several deep learning architectures such as convolutional deep neural networks (CNNs), and recurrent neural networks have been applied to computer vision, speech recognition, natural language processing, and audio recognition to produce state-of-the-art results on various tasks.

Amazon Rekognition is a part of the Amazon AI family of services. Amazon AI services use deep learning to understand images, turn text into lifelike speech, and build intuitive conversational text and speech interfaces.

Q: Do I need any deep learning expertise to use Amazon Rekognition?

A: No. With Amazon Rekognition, you don't have to build, maintain or upgrade deep learning pipelines.

To achieve accurate results on complex computer vision tasks such as object and scene detection, face analysis, and face recognition, deep learning systems need to be tuned properly and trained with massive amounts of labeled ground truth data. Sourcing, cleaning, and labeling data accurately is a time-consuming and expensive task. Moreover, training a deep neural network is computationally expensive and often requires custom hardware built using Graphics Processing Units (GPU).

Amazon Rekognition is fully managed and comes pre-trained for image and video recognition tasks, so that you don't have invest your time and resources on creating a deep learning pipeline. Amazon Rekognition continues to improve the accuracy of its models by building upon the latest research and sourcing new training data. This allows you to focus on high-value application design and development.

Webinar Series on Remote Learning

FAQs

Q: What are the most common use cases for Amazon Rekognition?

A: The most common use-cases for Rekognition Image include:

- Searchable Image Library
- Face-Based User Verification
- Sentiment Analysis
- Facial Recognition
- Image Moderation

The most common use-cases for Rekognition Video include:

- Search Index for video archives
- Easy filtering of video for explicit and suggestive content

Q: How do I get started with Amazon Rekognition?

A: If you are not already signed up for Amazon Rekognition, you can click the "Try Amazon Rekognition" button on the Amazon Rekognition page and complete the sign-up process. You must have an Amazon Web Services account; if you do not already have one, you will be prompted to create one during the sign-up process. Once you are signed up, try out Amazon Rekognition with your own images and videos using the Amazon Rekognition Management Console or download the Amazon Rekognition SDKs to start creating your own applications. Please refer to our step-by-step Getting Started Guide for more information.

Q: What APIs does Amazon Rekognition offer?

A: Amazon Rekognition Image offers APIs to detect objects and scenes, detect and analyze faces, recognize celebrities, detect inappropriate content, and search for similar faces in a collection of faces, along with APIs to manage resources. Rekognition Image also offers APIs to compare faces and extract text, while Rekognition Video also offers APIs to track persons and manage live stream video from Acuity. For details, please refer to the Amazon Rekognition API Reference.

FAQs

Q: What image and video formats does Amazon Rekognition support?

A: Amazon Rekognition Image currently supports the JPEG and PNG image formats. You can submit images either as an S3 object or as a byte array. Amazon Rekognition Video operations can analyze videos stored in Amazon S3 buckets. The video must be encoded using the H.264 codec. The supported file formats are MPEG-4 and MOV. A codec is software or hardware that compresses data for faster delivery and decompresses received data into its original form. The H.264 codec is commonly used for the recording, compression and distribution of video content. A video file format may contain one or more codecs. If your MOV or MPEG-4 format video file does not work with Rekognition Video, check that the codec used to encode the video is H.264.

Q: What file sizes can I use with Amazon Rekognition?

A: Amazon Rekognition Video supports up to 8 GB files and up to 2 hour videos when passed through as an S3 file.

Q: How does image resolution affect the quality of Rekognition Image API results?

A: Amazon Rekognition works across a wide range of image resolutions. For best results we recommend using VGA (640x480) resolution or higher. Going below QVGA (320x240) may increase the chances of missing faces, objects, or inappropriate content; although Amazon Rekognition accepts images that are at least 80 pixels in both dimensions.

Q: How small can an object be for Amazon Rekognition Image to detect and analyze it?

A: For example, if you are working with a 1600x900 image, the smallest face or object should be at least 45 pixels in either dimension.

FAQs

Amazon RDS

Q: What is Amazon RDS?

A: Amazon Relational Database Service (Amazon RDS) is a managed service that makes it easy to set up, operate, and scale a relational database in the cloud. It provides cost-efficient and resizable capacity, while managing time-consuming database administration tasks, freeing you up to focus on your applications and business.

Amazon RDS gives you access to the capabilities of a familiar MySQL, MariaDB, Oracle, SQL Server, or PostgreSQL database. This means that the code, applications, and tools you already use today with your existing databases should work seamlessly with Amazon RDS. Amazon RDS can automatically back up your database and keep your database software up to date with the latest version. You benefit from the flexibility of being able to easily scale the compute resources or storage capacity associated with your relational database instance. In addition, Amazon RDS makes it easy to use replication to enhance database availability, improve data durability, or scale beyond the capacity constraints of a single database instance for read-heavy database workloads. As with all Amazon Web Services, there are no up-front investments required, and you pay only for the resources you use.

Q: Which relational database engines does Amazon RDS support?

A: Amazon RDS supports Amazon Aurora, MySQL, MariaDB, Oracle, SQL Server, and PostgreSQL database engines.

Q: What does Amazon RDS manage on my behalf?

A: Amazon RDS manages the work involved in setting up a relational database: from provisioning the infrastructure capacity you request to installing the database software. Once your database is up and running, Amazon RDS automates common administrative tasks such as performing backups and patching the software that powers your database. With optional Multi-AZ deployments, Amazon RDS also manages synchronous data replication across Availability Zones with automatic failover.

Since Amazon RDS provides native database access, you interact with the relational database software as you normally would. This means you're still responsible for managing the database settings that are specific to your application. You'll need to build the relational schema that best fits your use case and are responsible for any performance tuning to optimize your database for your application's workflow.

FAQs

Q: How can I get Amazon Rekognition predictions reviewed by humans?

A: Amazon Rekognition is directly integrated with Amazon Augmented AI (Amazon A2I) so you can easily route low confidence predictions from Amazon Rekognition Image to human reviewers. Using the Amazon Rekognition API for content moderation or the Amazon A2I console, you can specify the conditions under which Amazon A2I routes predictions to reviewers, which can be either a confidence threshold or a random sampling percentage. If you specify a confidence threshold, Amazon A2I routes only those predictions that fall below the threshold for human review. You can adjust these thresholds at any time to achieve the right balance between accuracy and cost-effectiveness. Alternatively, if you specify a sampling percentage, Amazon A2I routes a random sample of the predictions for human review. This can help you implement audits to monitor the prediction accuracy regularly. Amazon A2I also provides reviewers with a web interface consisting of all the instructions and tools they need to complete their review tasks. For more information about implementing human review with Amazon Rekognition, see the Amazon A2I webpage.

Q: How does video resolution affect the quality of Rekognition Video API results?

A: The system is trained to recognize faces larger than 32 pixels (on the shortest dimension), which translate into a minimum size for a face to be recognized that varies from approximately 1/7 of the screen smaller dimension at QVGA resolution to 1/30 at HD 1080p resolution. For example, at VGA resolution, users should expect lower performances for faces smaller than 1/10 of the screen smaller dimension.

Q: What else can affect the quality of the Rekognition Video APIs?

A: Besides video resolution, heavy blur, fast moving persons, lighting conditions, pose may affect the quality of the APIs.

Q: What is the preferred user video content that is suitable for Rekognition Video APIs?

A: This API works best with consumer and professional videos taken with frontal field of view in normal color and lighting conditions. This API is not tested for black and white, IR or extreme lighting condition. Applications that are sensitive to false alarms are advised to discard outputs with confidence score below a selected (application-specific) confidence score.

Q: In which AWS regions is Amazon Rekognition available?

A: For a list of all regions where Amazon Rekognition is available, see the AWS Region table.

RESOURCES

[Amazon Rekognition](#) - Amazon Rekognition provides fast and accurate face search, allowing you to identify a person in a photo or video using your private repository of face images.

[Amazon Rekognition User Guide](#) – Getting Started with Amazon Rekognition

[AWS Educate](#) – AWS Educate gives students and educators access to content and programs that enable them to skill up for cloud careers in growing fields. AWS Educate also connects companies hiring for cloud skills to qualified student job seekers with the AWS Educate Job Board.

[AWS Educate Office Hours for Educators and Students](#) – Webinars, office hours, and training sessions for educators and students

[AWS User Groups](#) - User groups are peer-to-peer communities which meet regularly to share ideas, answer questions, and learn about new services and best practices.