
AWS
re:Invent

AIM318

Which episode? Using AI to streamline media content operations

Anushri Mainthia

Sr. Product Manager,
Amazon Rekognition
AWS

Sally Hubbard

Director, Media Management
Public Broadcasting Service (PBS)

AWS re:Invent

© 2020, Amazon Web Services, Inc. or its affiliates. All rights reserved.

What to expect

Media trends

Introduction to Amazon Rekognition

Use cases deep dive

PBS and *Sesame Street*

Media trends



Explosion of direct to consumer content / platforms



Accelerating demand for content



Enhanced audience experiences

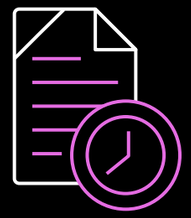


Monetization innovation

Viewers expect more

Media customers need to accelerate innovation to differentiate themselves and maintain profitability

Challenges



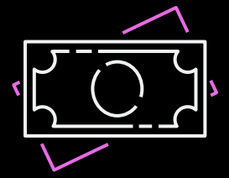
Inefficient, manual, and tedious



Subjective leading to inconsistencies (‘Blink and you miss it’)



Missed opportunities



Expensive



Not scalable

Amazon Rekognition image and video



OBJECT, SCENE, AND ACTIVITY



NEW! CUSTOM LABELS



UPDATED! CONTENT MODERATION



TEXT



PPE DETECTION



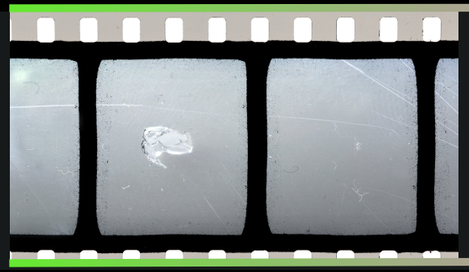
FACE DETECTION AND ANALYSIS



FACE COMPARE & SEARCH



CELEBRITY RECOGNITION



NEW! VIDEO SEGMENTS



LIVE STREAM VIDEO



PATHING

Media use cases

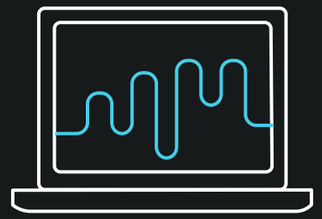


DISCOVERY

Rich asset search

Engaging viewer experiences

Content creation – highlights and promos



OPERATIONS

Quality control

Monitoring and compliance

Post production workflows



MONETIZATION

Ad operations

Native and contextual advertising

Reporting and analytics

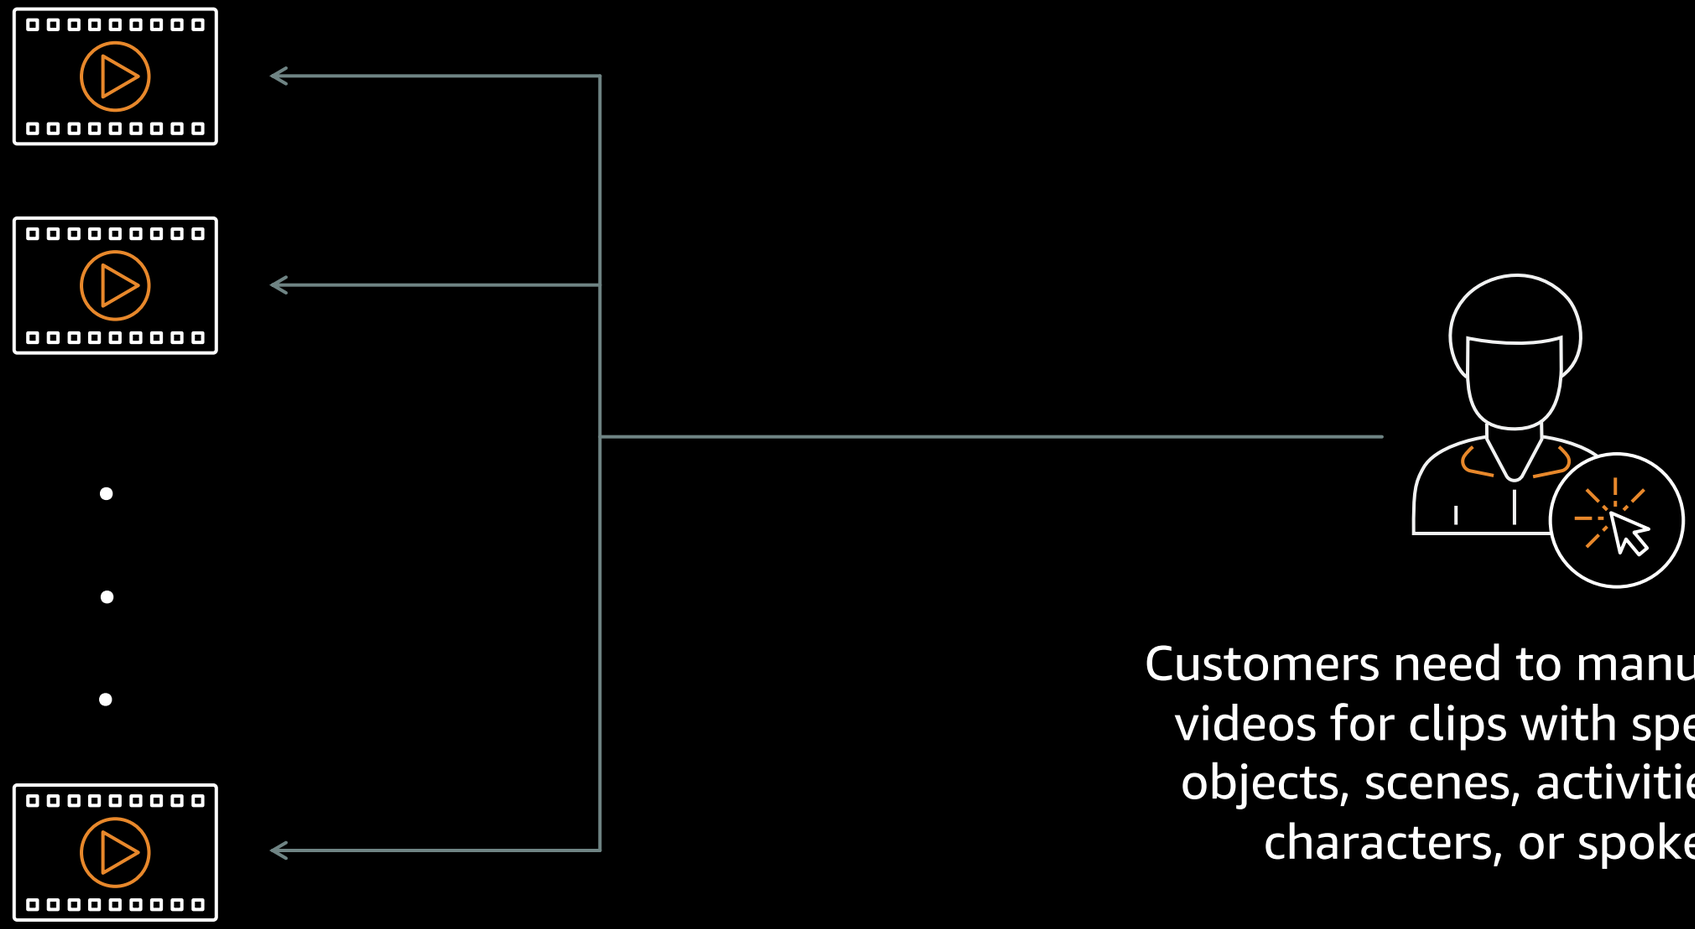


Discovery

© 2020, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Discovery | Operations

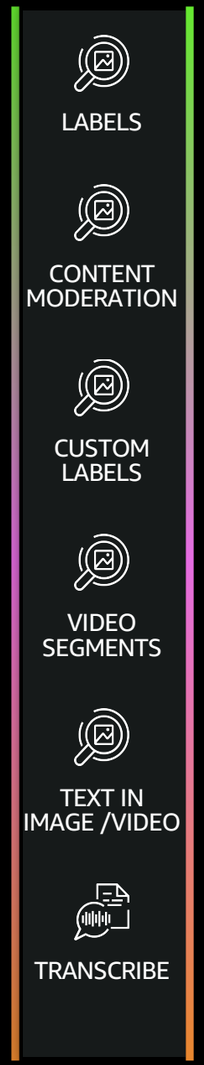
The challenge

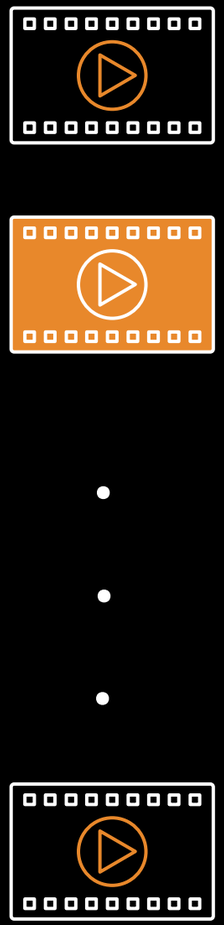


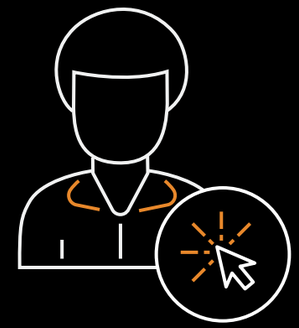
Customers need to manually search all videos for clips with specific people, objects, scenes, activities, animated characters, or spoken words

Discovery | Operations

Using Amazon Rekognition







Easily find specific videos and clips within your library using rich metadata from AI services

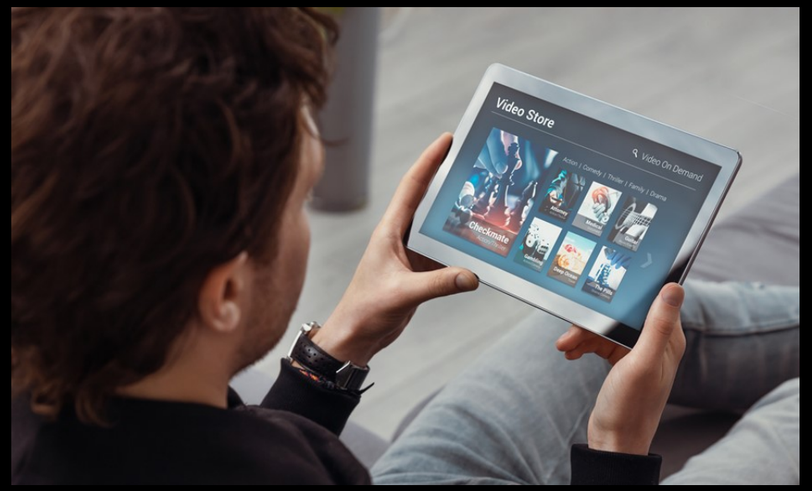
"All "All videos with Jeff Bezos" "All videos with Jeff Bezos and Werner Vogels"

Discovery use cases

Discovery | Operations



Automatic creation of rich and relevant metadata



Enhanced viewer experiences



Faster content creation



PBS, *Sesame Street* and creating content for America's largest classroom

Sally Hubbard

Director, Media Management
Public Broadcasting Service (PBS)

Who we are

PBS is a private, nonprofit corporation—founded in 1969—whose members are America's public TV stations, serving all 50 states, Puerto Rico, U.S. Virgin Islands, Guam, and American Samoa

America's largest and most trusted classroom



approximately

7 PB/60,000 titles

in the PBS digital library



MASTERPIECE

Millions

of moments

The challenge

How do we go from traditional TV metadata and episode level information to *moments?*



SEASON 41, EPISODE 36:
MINE-ITIS ON SESAME STREET

First aired: December 17, 2010

Elmo and Leela investigate when a Brian Williams report indicates that no one is sharing on Sesame Street. They discover that, indeed, the disease of "mine-itis" is running rampant.

The use case

Who's on *Sesame Street* … what are they doing, and why is that important?

Sesame Street

PBS Enterprise Knowledge Graph

How to make them meet?



Topics and concepts

SHARING
GIVING
TAKING TURNS

Linking metadata to the PBS taxonomy tied to business needs



[Brian Williams]



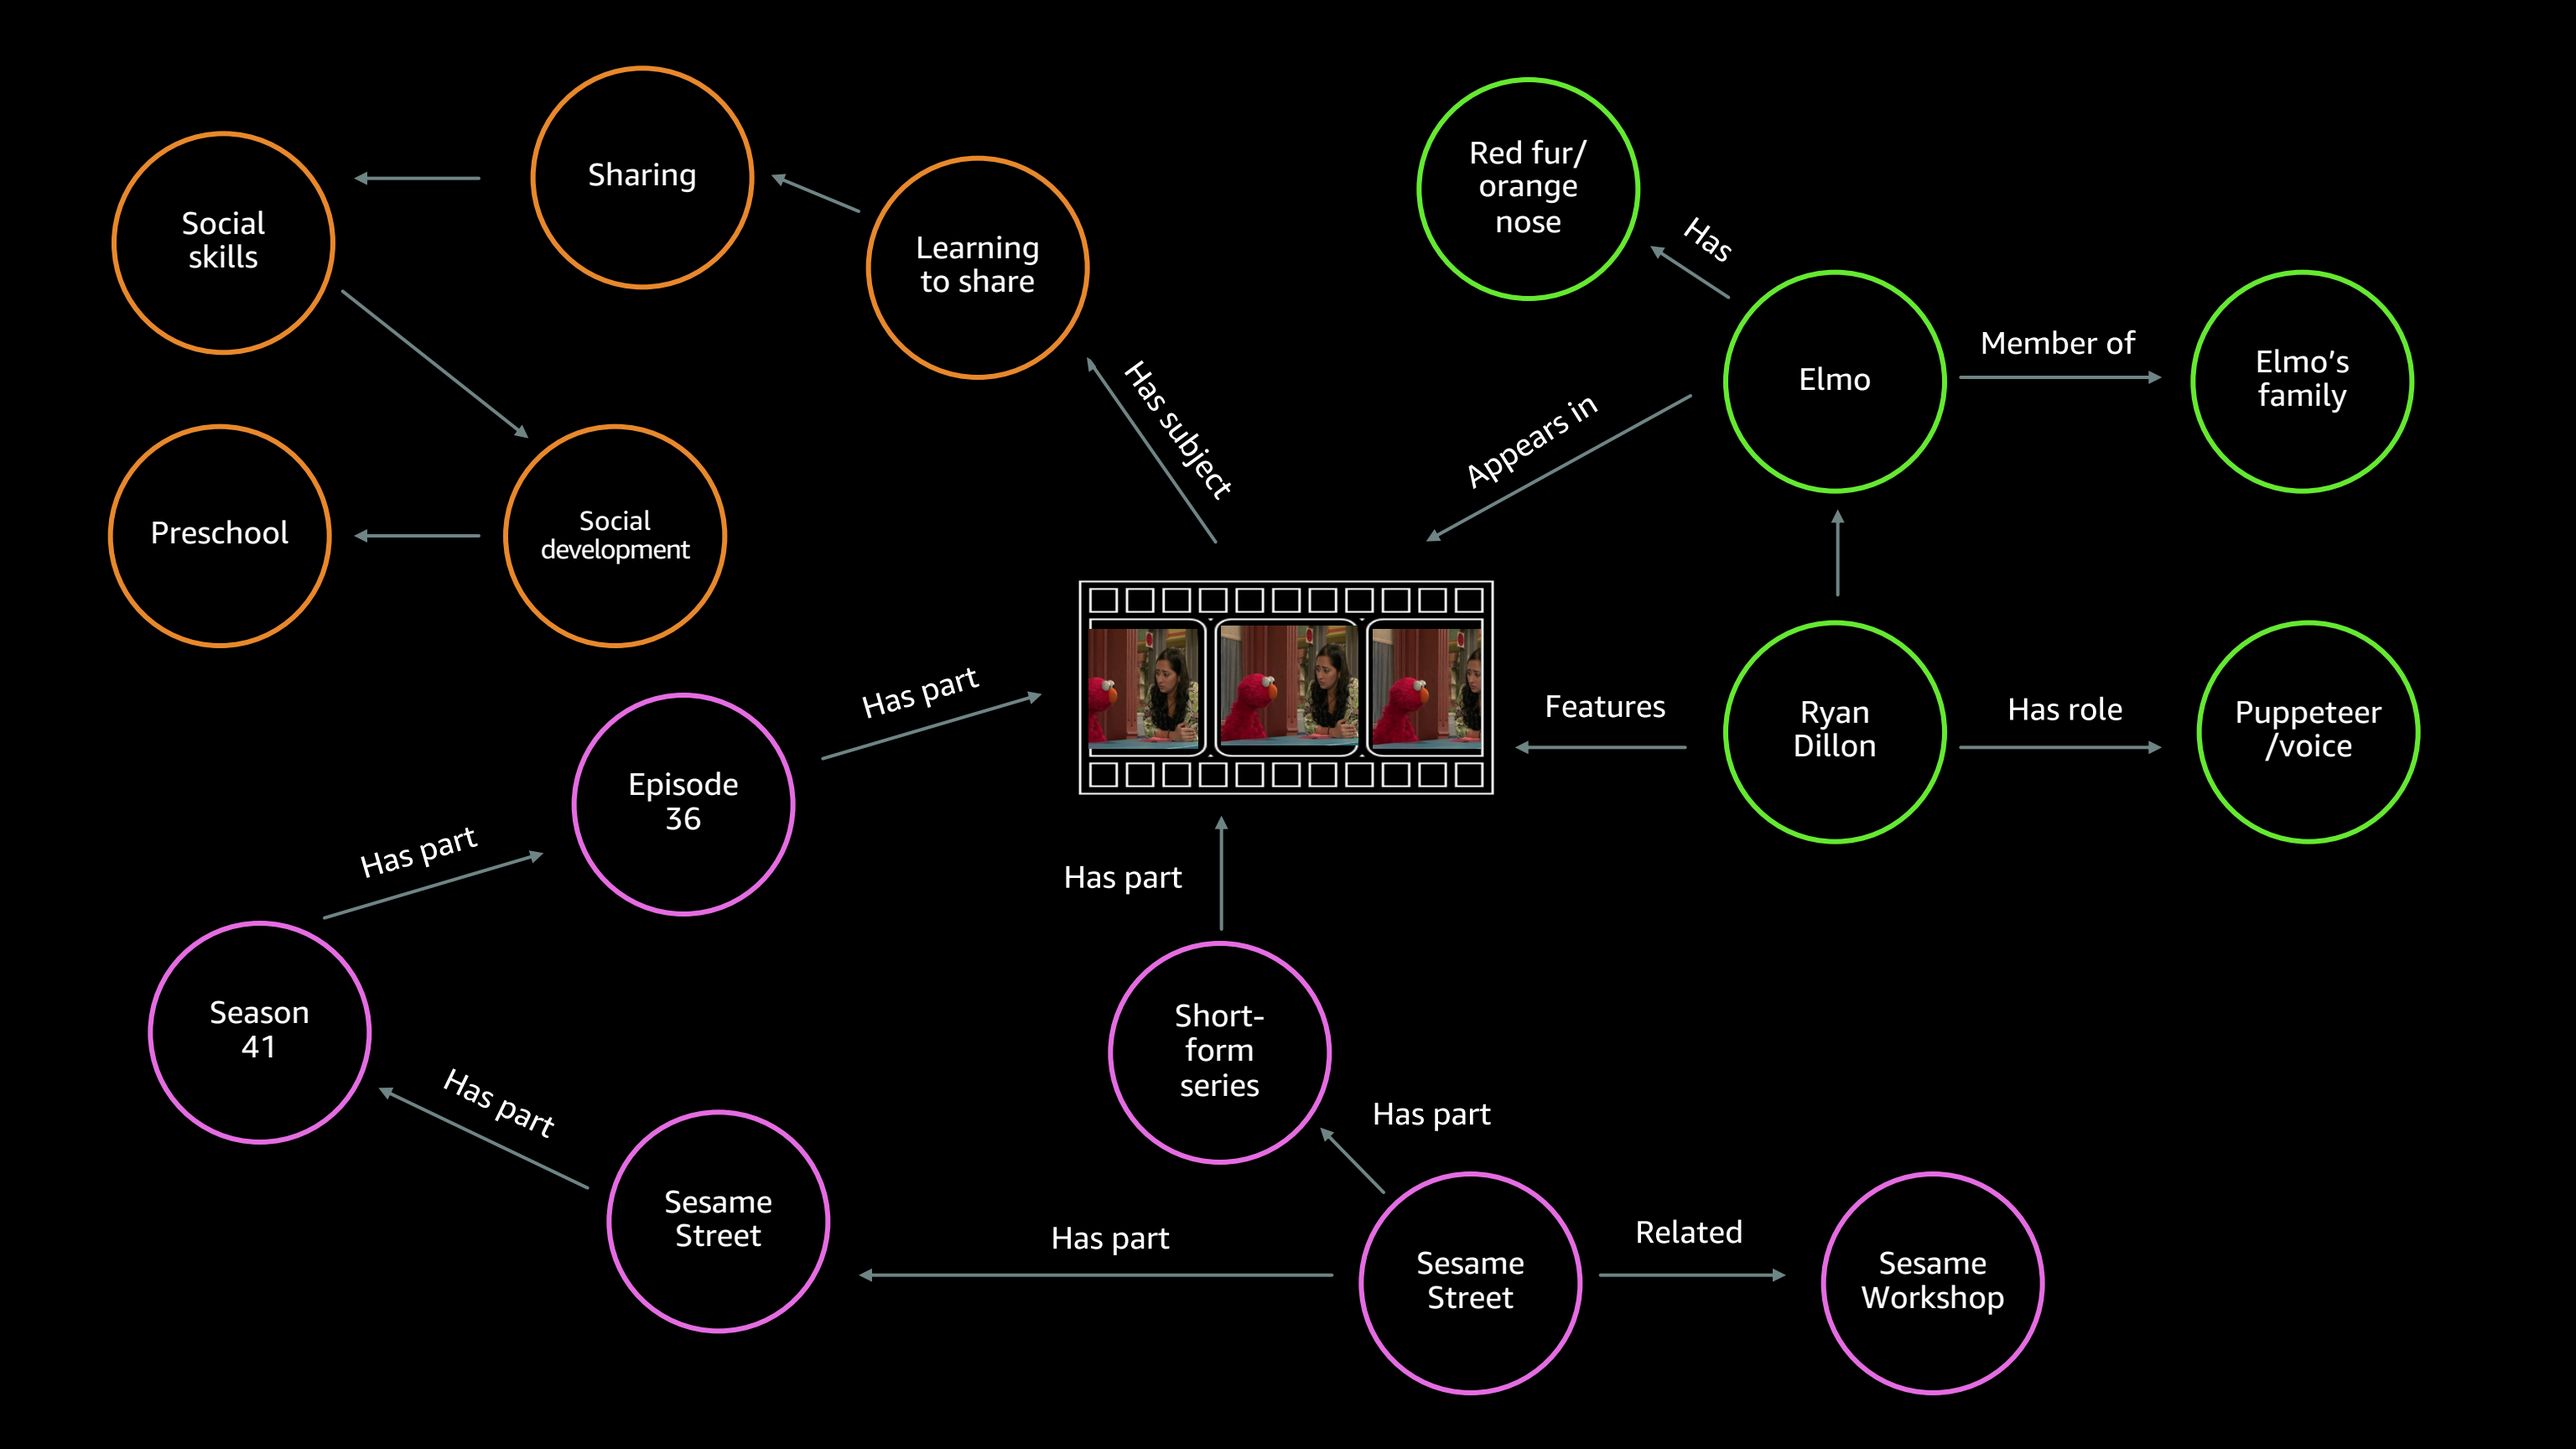
"Microphone"
"Sign"

[Nitya Vidyasagar]





Muppet [Elmo]
Surprise (99.93%)
Listening (15.14%)
Learning (8.65%)


Social skills
Sharing
Learning to share
Has subject
Red fur/ orange nose
Has
Elmo
Member of
Elmo's family
Appears in
Preschool
Social development
Has part
Episode 36
Features
Ryan Dillon
Has role
Puppeteer /voice
Has part
Season 41
Has part
Has part
Short-form series
Has part
Sesame Street
Has part
Sesame Street
Related
Sesame Workshop


Does machine learning work with muppets?

The solution

PBS

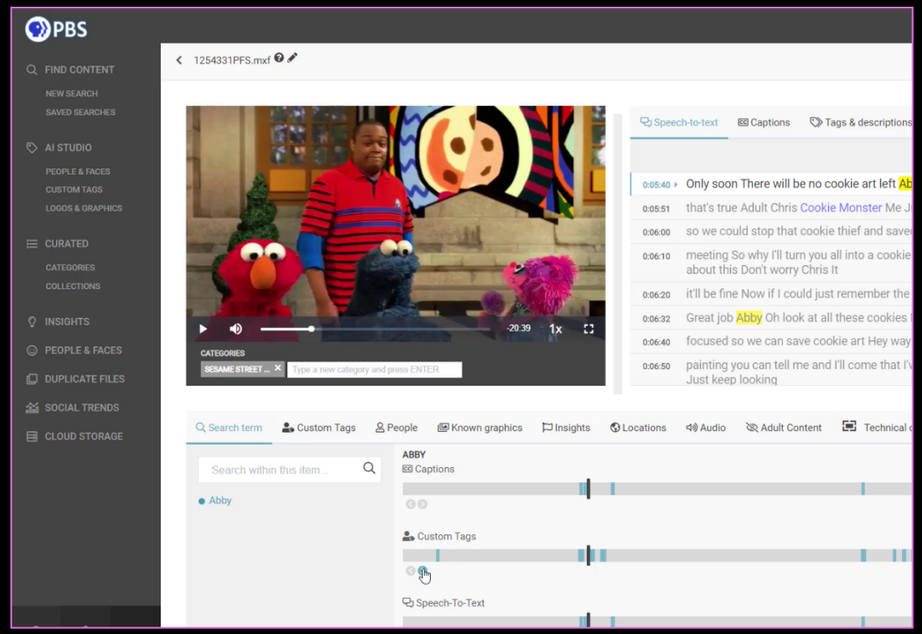
MEDIA AND CONTENT MANAGEMENT

Content upload to S3

Knowledge Graph (PoolParty)

GRAYMETA

CURIO ML PLATFORM USER INTERFACE



aws

AMAZON REKOGNITION

- Labels
- Video segment detection
- Custom labeling and tags
- Celebrity recognition
- Content moderation
- Text detection

The results

High-quality, rich, and consistent metadata

We have found a scalable mechanism to provide new, rich moment-based metadata to complement existing data stores

Search and discovery

ML can be combined with a knowledge model and integrated into recommendations and delivery

Find new uses

We can find new value in old content, and may also support creation of new content

Cost and time efficiency

ML processing has potential to not only improve existing processes, but to make entirely new processes possible

What’s next

With this production-ready solution to automatically generate rich, relevant, and customized metadata (in a fraction of the time required to do so manually), we see many possibilities to surface new content from our archive, create engaging user experiences, and more

Scale

Expand to other PBS shows and programming to drive overall search and discovery

Recommendations

Use metadata and taxonomy to surface relevant content recommendations to drive viewer experience

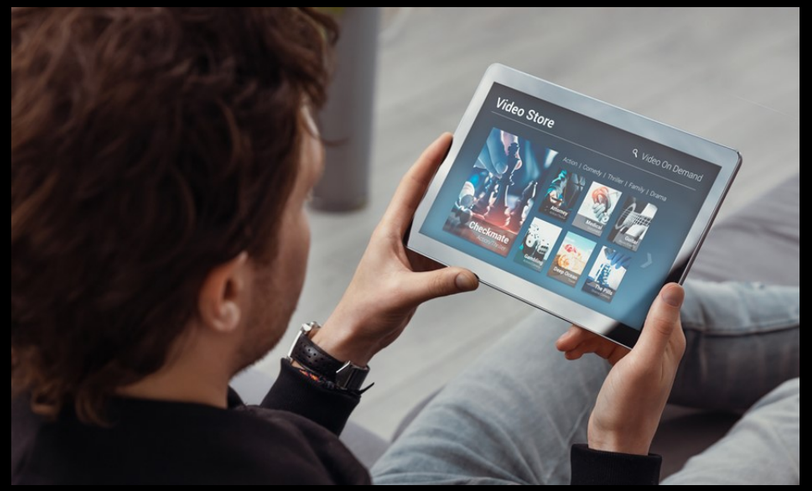


Discovery use cases

Discovery | Operations



Automatic creation of rich and relevant metadata



Enhanced viewer experiences



Faster content creation

GrayMeta Curio Platform

Discovery | Operations

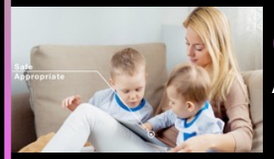
UNDER THE HOOD



CUSTOM LABELS
Detect custom objects like characters, logos, or products, as well as custom scene and concepts such as shot types (e.g., wide shot, over-the-shoulder shot), emotions (romance, sadness)



LABELS: OBJECTS, SCENES, AND ACTIVITIES
Leverage deep learning models trained on thousands of objects, scenes, and activities



CONTENT MODERATION
Automatically detect inappropriate or unwanted content



CELEBRITY RECOGNITION
Leverage deep learning models trained on hundreds of thousands of celebrities



VIDEO SEGMENTS
Automate detection of shot boundaries, black frames, color bars, and end credits



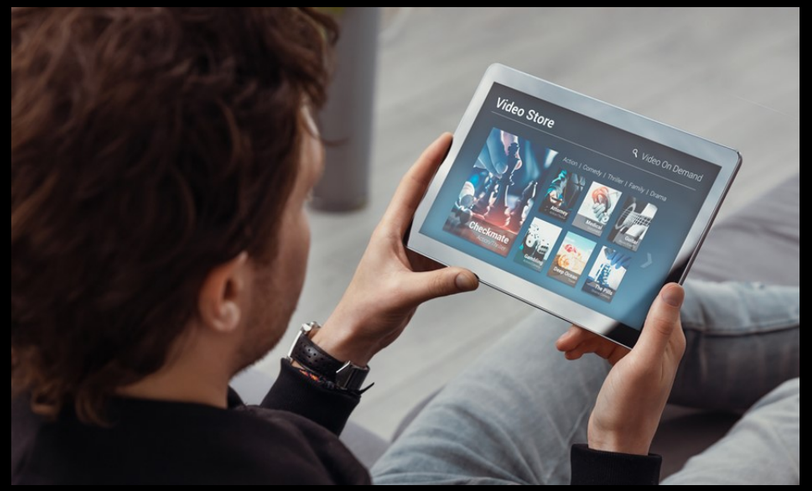
TEXT IN IMAGE AND VIDEO
Detect text in content, slates, and more

Discovery use cases

Discovery | Operations



Automatic creation of rich and relevant metadata



Enhanced viewer experiences



Faster content creation

Enhanced viewer experiences

Discovery | Operations

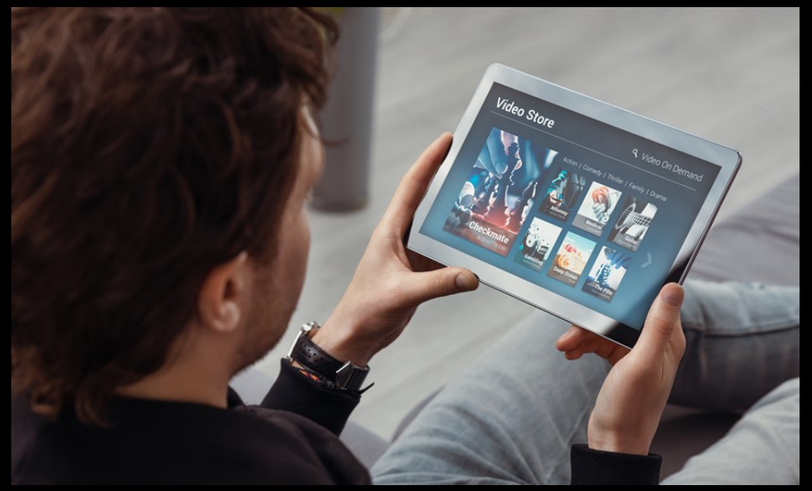


Discovery use cases

Discovery | Operations



Automatic creation of rich and relevant metadata



Enhanced viewer experiences



Faster content creation

Case study

Discovery | Operations



PROBLEM STATEMENT
Can we automate and scale promo creation

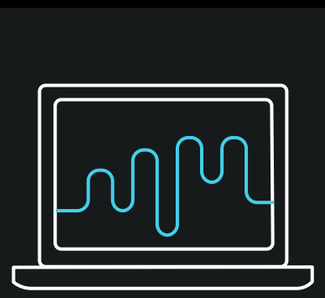
CHALLENGE
A 30" rough edit requires at least 5 days of manual viewing and logging. To get scale, customers often leverage external agencies and have to consolidate results. *Process is onerous, timing consuming, and inefficient.*

SOLUTION
PromoMii launched Nova, an automated editing tool, that leverages Amazon Rekognition – now, creating a clip takes less than 2 hours vs. 2 days

The results



Blog and reference architecture



Media operations

AWS re:Invent

© 2020, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Media operations use cases

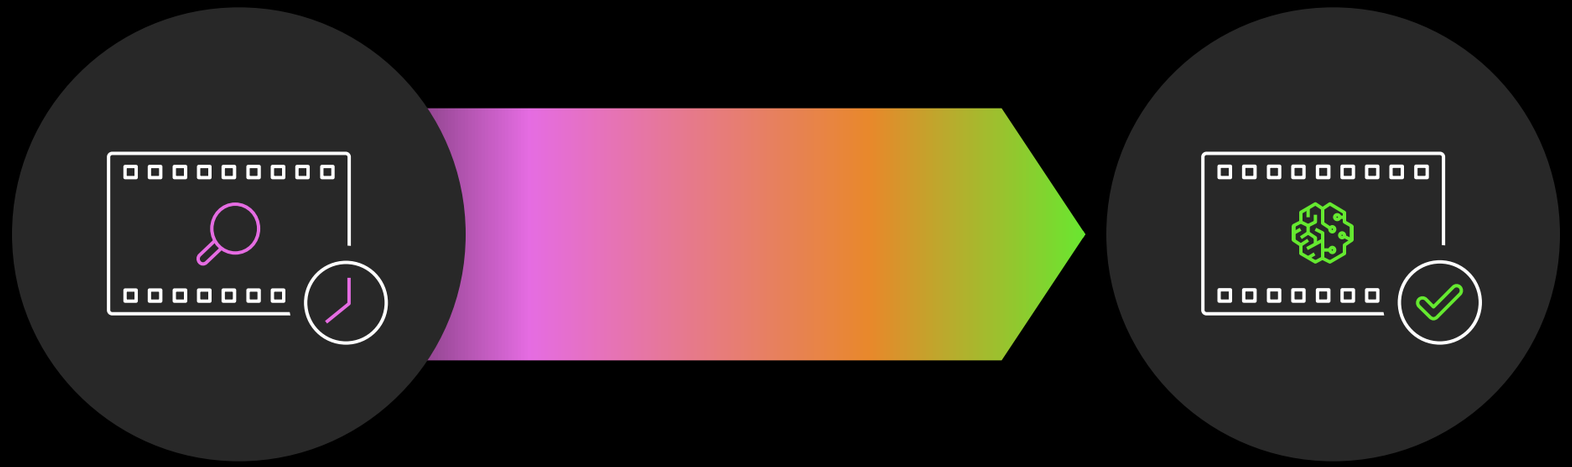


Streaming	Monitoring broadcast feeds
File-based workflows	Reviewing content for quality control and compliance

Image licensed under CC By 2.0; retrieved from https://en.wikipedia.org/wiki/Master_control#/media/File:111807y.jpg

Streamlining operations

FUTURE VISION

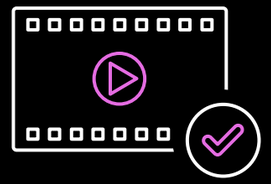
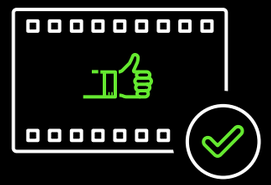
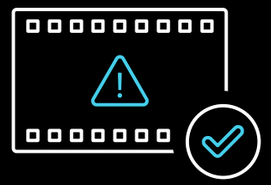


Tedious, manual review of every frame

Automated content monitoring with AI/ML

Only review content that is flagged by ML, and rest of content is 'assumed good'

Operational workflows

	Rekognition APIs
 **Content preparation** Inserting binge markers such as ‘Next Episode,’ removing color bars and slates, detecting commercial black frames	Video segment detection, Custom Labels, Labels, Text detection
 **Quality control** Correctness of program and station logos, start and end of program, accidental text, etc.	Video segment detection, Custom Labels, Content moderation, Text detection, Labels
 **Compliance** Check content against global standards and practices guidelines	Video segment detection, Content moderation, Custom Labels, Text detection, Labels

Discovery | **Operations**

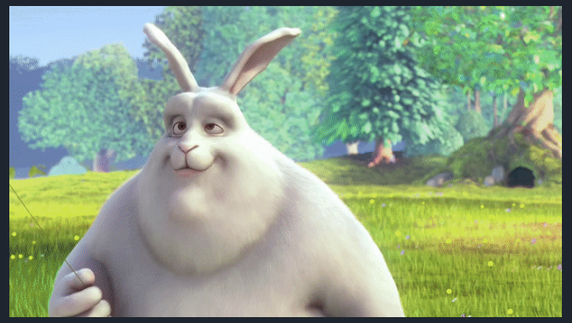
Using Rekognition

UNDER THE HOOD – VIDEO SEGMENT DETECTION

Black frames detection



End credits detection



Shot detection



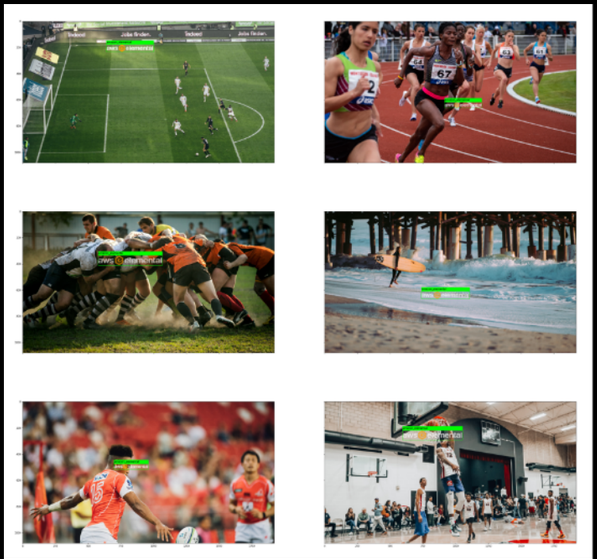
Color bars

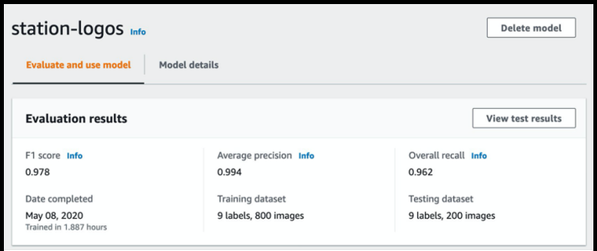


Using Rekognition

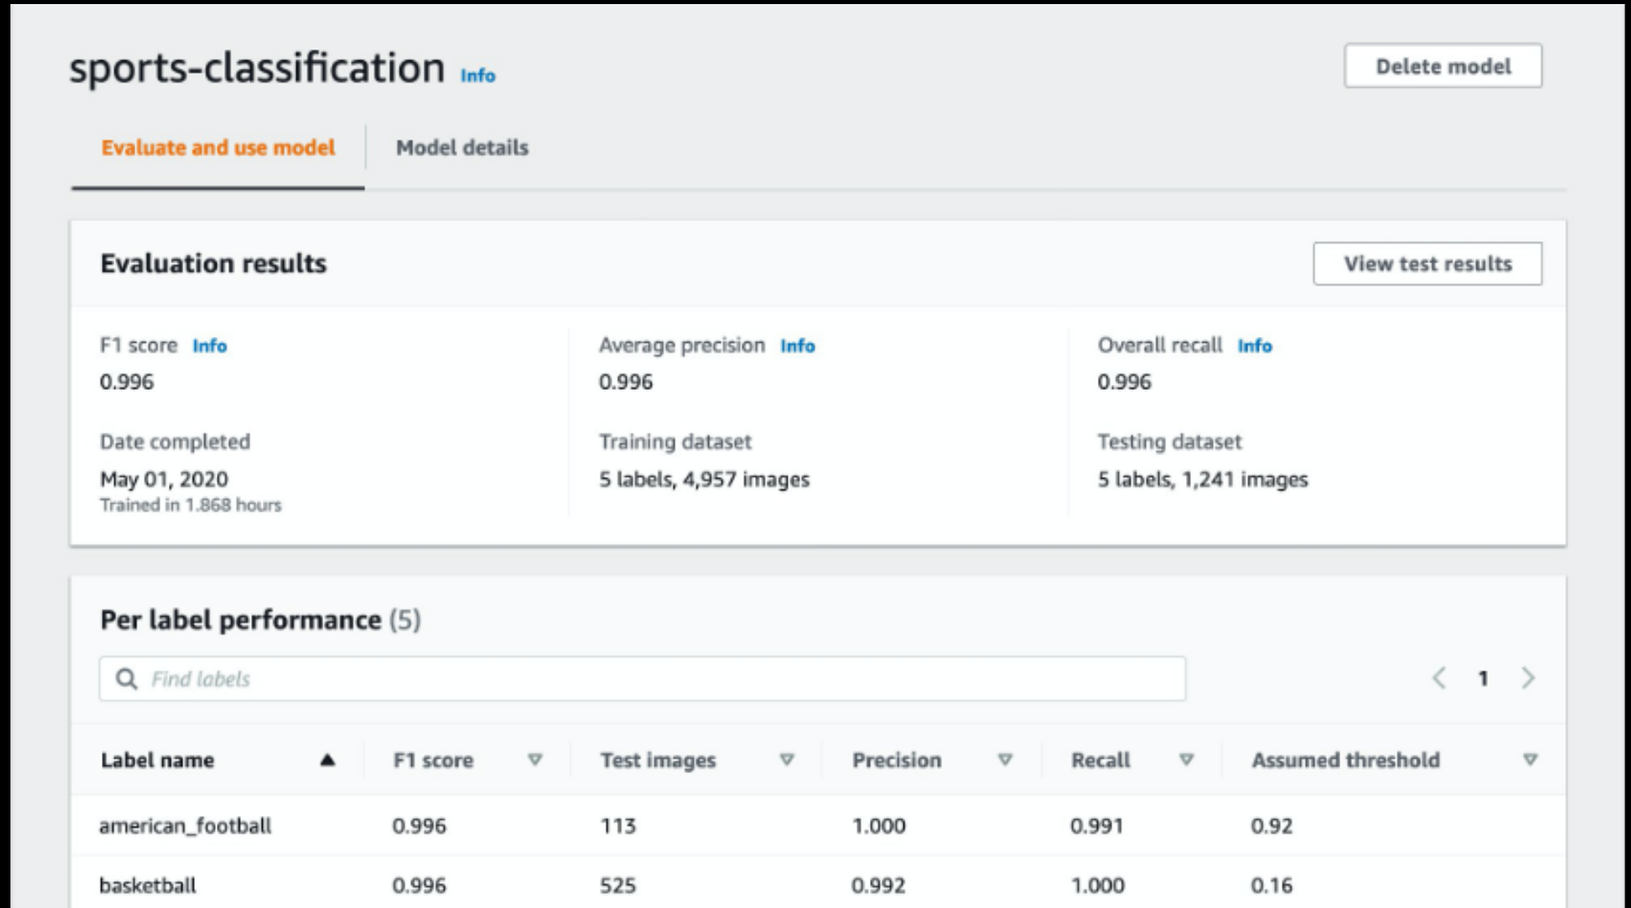
UNDER THE HOOD – CUSTOM LABELS

Station logo monitoring





Sports classifier



Using Rekognition

UNDER THE HOOD – CONTENT MODERATION

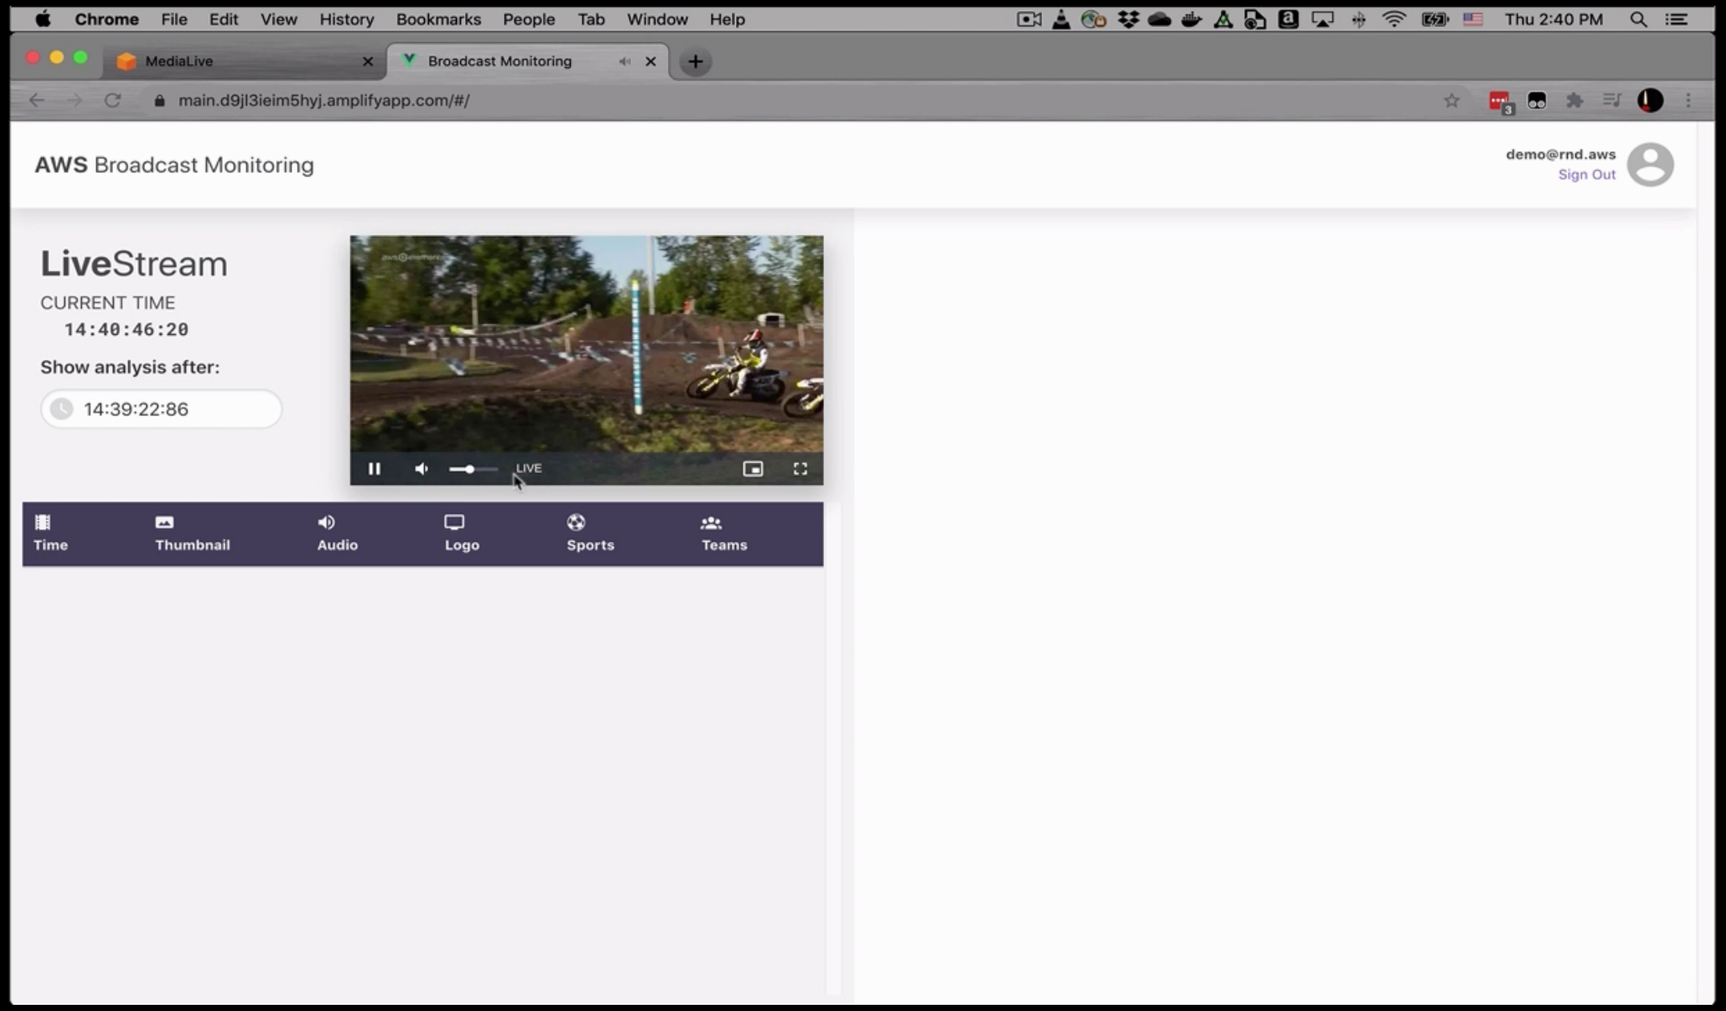
Top-level category	Second-level category
Explicit Nudity	Nudity Graphic Male Nudity Graphic Female Nudity Sexual Activity Illustrated Nudity Adult Toys
Suggestive	Female Swimwear or Underwear Male Swimwear or Underwear Partial Nudity Revealing Clothes
Violence	Graphic Violence or Gore Physical Violence Weapon Violence Weapons Self-Injury
Visually Disturbing	Emaciated Bodies Corpses Hanging

Top-level category	Second-level category
Rude Gestures	Middle Finger
Drugs	Drug Products Drug Use Pills Drug Paraphernalia
Tobacco	Tobacco Products Smoking
Alcohol	Alcoholic Beverages Drinking
Gambling	Gambling
Hate Symbols	Nazi Party White Supremacy Extremist

Amazon Rekognition does not flag content as “safe” or “unsafe”

Instead, you get granular output about the nature of the content. You can apply your own policies that define what is safe for your customers.

Operations solution demo



Blog and reference architecture

Media attribution: *Motocross & Soccer Footage: Media assets courtesy of AWS Elemental || Football clip Footage: Northwestern vs. Michigan | Fox College Football Highlights | Licensed under CC (http://creativecommons.org/licenses/by/3.0) | Source: https://youtube.com/watch?v=vQ7n7nGcTWE*

Benefits of Amazon Rekognition for media

Increased efficiency

Automate metadata tagging with Rekognition to reduce manual content review and enable teams to focus on higher-value tasks

Highly scalable and highly available

No machine learning expertise required

Leverage deep learning-based image and video analysis or create your own CV model using your own data without any ML expertise

Easy to use

Fully managed APIs, create custom inference APIs with little to no coding required



What’s next?

© 2020, Amazon Web Services, Inc. or its affiliates. All rights reserved.

How to get started

PARTNERS

DEVELOPER GUIDE

VIDEOS AND TUTORIAL



More resources

© 2020, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Resources

AMAZON REKOGNITION BLOGS AND SOLUTIONS



Synchronized and TF1 use Amazon Rekognition to automate thumbnail creation
https://aws.amazon.com/fr/solutions/case-studies/customer-references-tf1-synchronized/



ESSENCE as Starchive and AWS help deliver the 2020 ESSENCE Festival of Culture
https://aws.amazon.com/blogs/media/successful-innovation-for-essence-as-starchive-and-aws-help-deliver-the-2020-essence-festival-of-culture/



PromoMii: Video ads powered by AWS machine learning
https://aws.amazon.com/blogs/startups/promomii-video-ads-powered-by-aws-machine-learning/



VideoFashion: Using Amazon Rekognition to make your video archive fashionable
https://www.youtube.com/watch?v=mSsT95SDqzg&feature=emb_logo

Resources

AMAZON REKOGNITION BLOGS AND SOLUTIONS



Media2Cloud solution: A serverless framework for media archive ingest to the cloud
https://aws.amazon.com/blogs/media/introducing-media2cloud-a-serverless-framework-for-media-archive-ingest-to-the-cloud/



AWS Media Insights Engine
https://github.com/awslabs/aws-media-insights-engine/



Automating broadcast video monitoring with Amazon Rekognition
https://aws.amazon.com/blogs/media/automate-broadcast-video-monitoring-using-machine-learning-on-aws/



Inserting ad breaks into video content using Amazon Rekognition
https://aws.amazon.com/blogs/media/inserting-ad-breaks-into-video-content-using-amazon-rekognition-aws-elemental-mediaconvert-and-aws-elemental-mediatailor/



Custom Brand Detection Solution using Amazon Rekognition Custom Labels
https://github.com/aws-samples/amazon-rekognition-custom-brand-detection

Resources

AMAZON REKOGNITION BLOGS AND SOLUTIONS



Amazon Rekognition Segment Detection
https://aws.amazon.com/blogs/media/streamline-media-analysis-tasks-with-amazon-rekognition-video/



Amazon Rekognition shot detection demo using Segment API
https://github.com/aws-samples/amazon-rekognition-shot-detection-demo-using-segment-api



Amazon Rekognition Custom Labels
https://aws.amazon.com/blogs/machine-learning/announcing-amazon-rekognition-custom-labels/



Using Amazon Rekognition Custom Labels with Amazon A2I for human review
https://aws.amazon.com/blogs/machine-learning/using-amazon-rekognition-custom-labels-and-amazon-a2i-for-detecting-pizza-slices-and-augmenting-predictions/

Thank you!



AWS re:Invent

© 2020, Amazon Web Services, Inc. or its affiliates. All rights reserved.



Please complete the session survey

AWS re:Invent

© 2020, Amazon Web Services, Inc. or its affiliates. All rights reserved.