



SUMMIT
ONLINE

INT 04

The effortless development of custom computer vision models

Sara van de Moosdijk

ML Partner Solutions Architect
Amazon Web Services



Building a model for Auslan

300

Estimated sign
languages worldwide

19th

Century origin of Australian
sign language (Auslan)






Video classification

Action recognition

Pose segmentation

Sequence to sequence



Building our own in Amazon SageMaker

- Default PyTorch environment
- Transfer learning on Resnet18
- 4 months
- 148 training jobs
- 3 hyper parameter tuning jobs
- 49 training hours
- Achieved 83% accuracy



Amazon SageMaker demo

Amazon Rekognition

Image and video analysis



Labels (object, scenes, and activities)



Unsafe image and video detection



Text in image



Pathing



Face search



Face detection and analysis



Celebrity recognition



Real-time video analysis

Amazon Rekognition Custom Labels

Customised image analysis to easily detect objects and scenes you define as most relevant to your domain



Guided experience
to create labeled
images



Train and evaluate
with **no coding** and
no **ML experience**



Easy-to-use
fully managed API

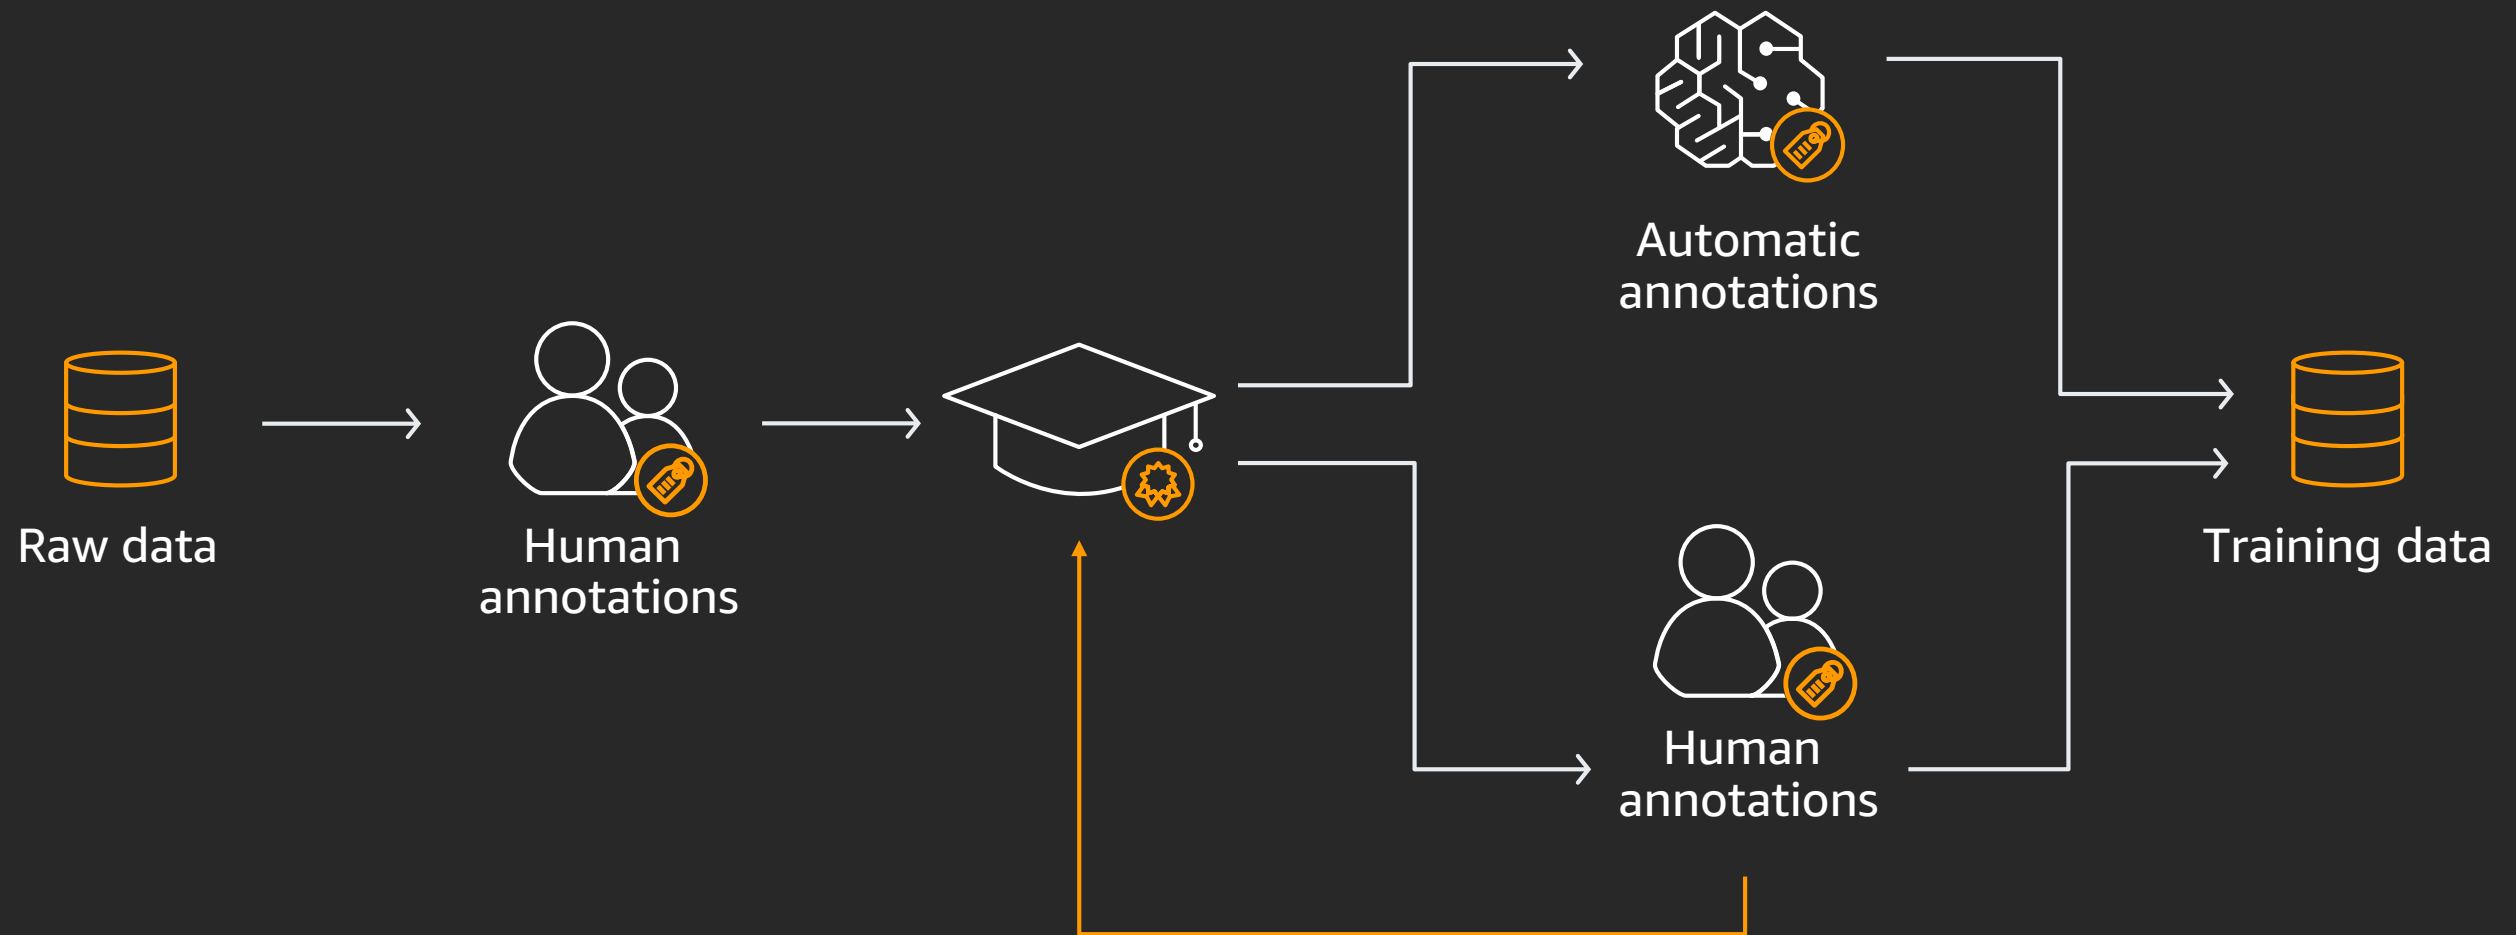
Amazon Rekognition demo

Creating datasets

- Import from Amazon S3 bucket
- Copy an existing Amazon Rekognition dataset
- Upload images from computer
- Amazon SageMaker Ground Truth

Amazon SageMaker Ground Truth

Build highly accurate training datasets using machine learning



Manifest file


```
{  
  "source-ref": "S3 bucket location",  
  "signs": 0,  
  "signs-metadata": {  
    "class-name": "hello",  
    "confidence": 1,  
    "type": "groundtruth/image-classification",  
    "human-annotated": "yes",  
    "creation-date": "2018-10-18T22:18:13.527256"  
  }  
}
```

Evaluating the model

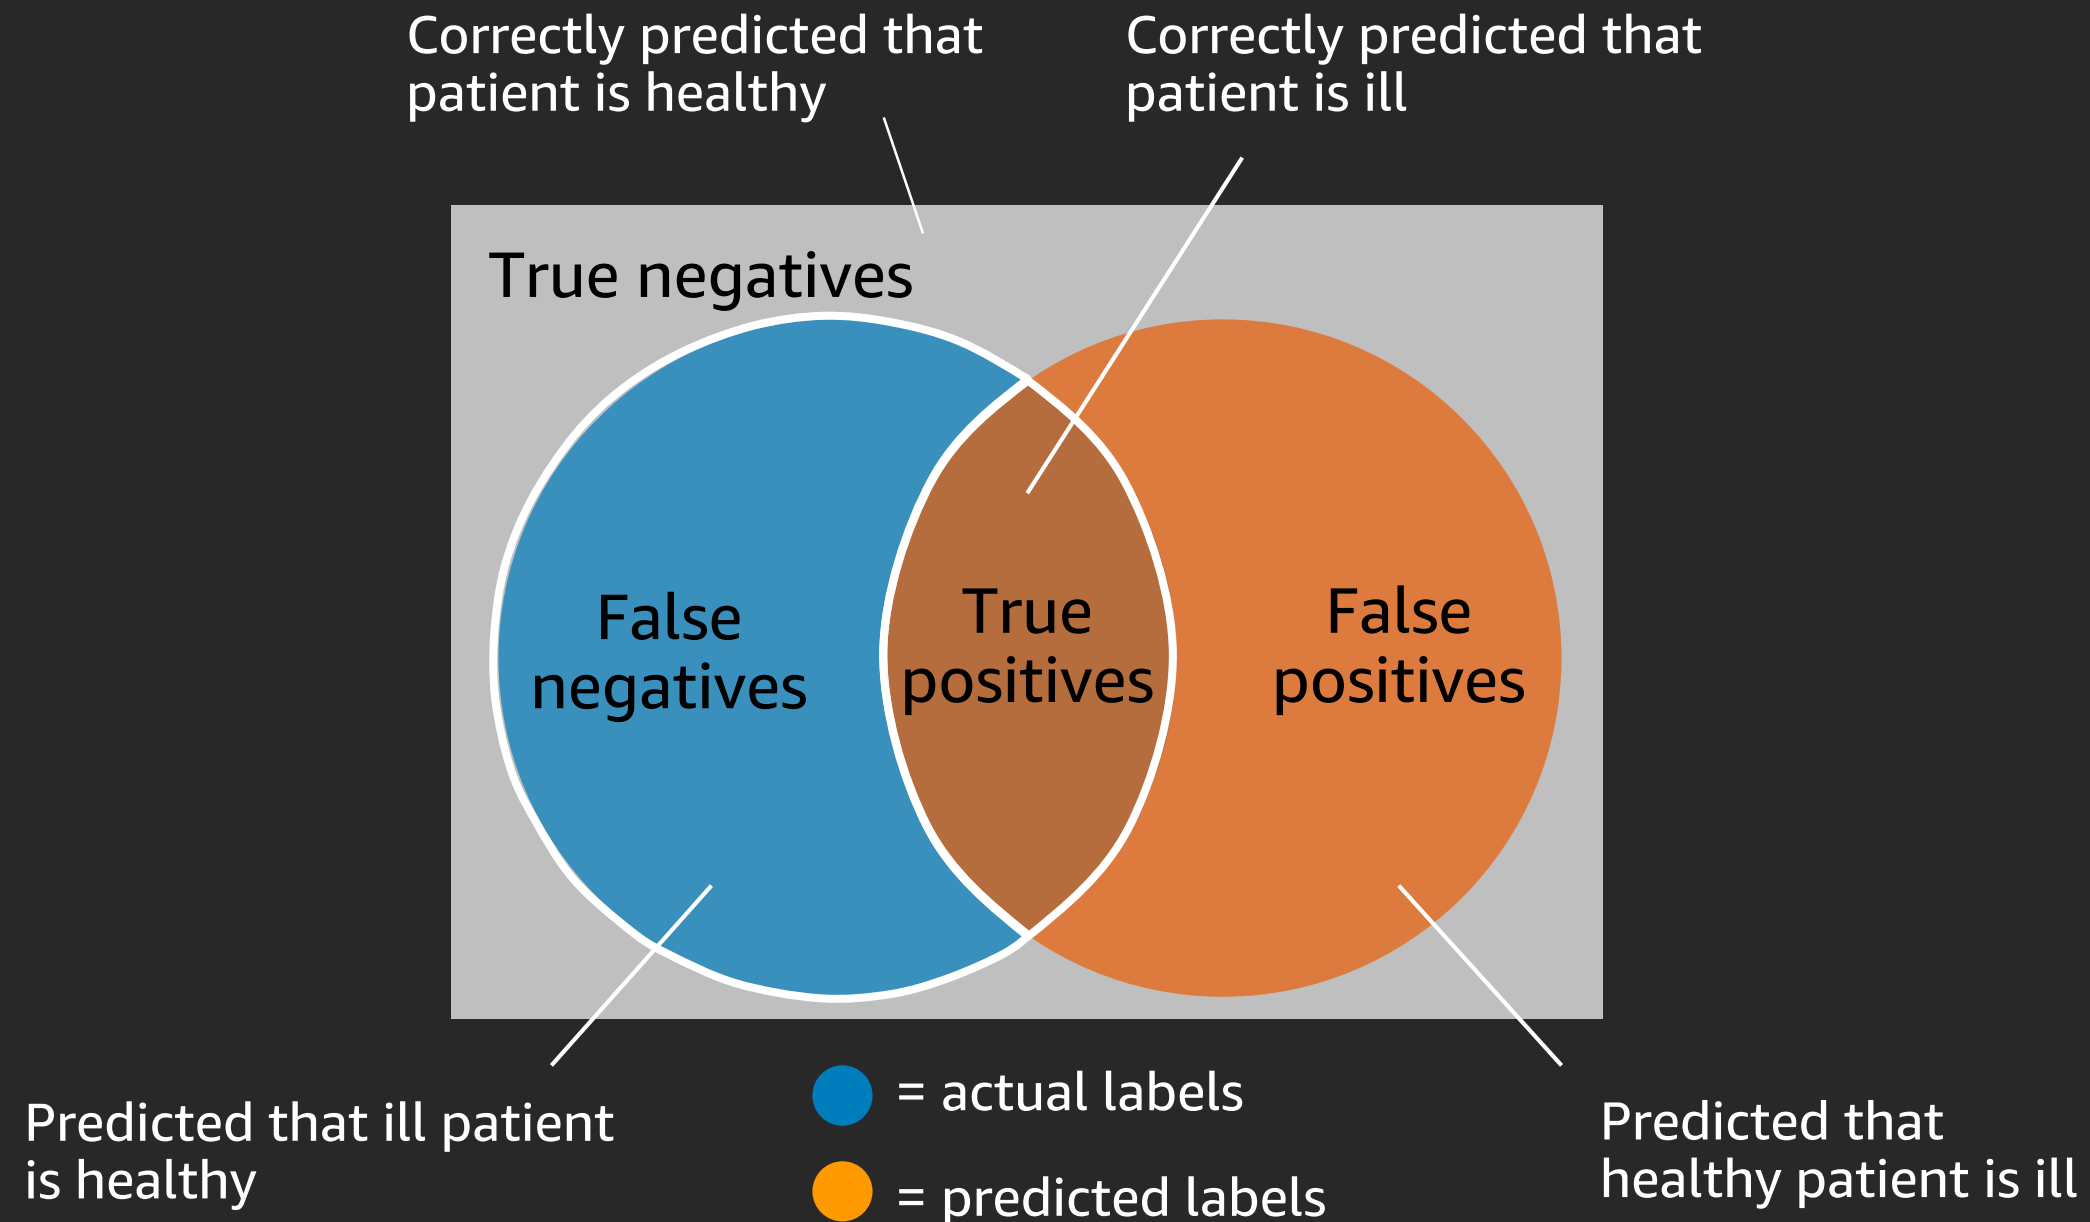
Console metrics

- Precision, recall, F1 score
- Mean average precision, mean average recall

Precision, recall, and F1 score



Precision, recall, and F1 score



Evaluating the model

Console metrics

- Precision, recall, F1 score
- Mean average precision, mean average recall

Additional API metrics

- Summary file
- Manifest snapshot

Evaluating the model

Console metrics

- Precision, recall, F1 score
- Mean average precision, mean average recall

Additional API metrics

- Summary file
- Manifest snapshot

Model feedback solution to improve model



Deploying the model

Start and stop the model endpoint

Assign inference units

- Priced at inference unit per hour
- Supports 5 transactions per second (TPS)

Provide image as base64-encoded byte array or as S3 object


Adjust threshold per request

Choosing the right tool



Amazon SageMaker

- Supports advanced ML use cases
- Full control of environment and implementation
- Available in SYD region



Amazon Rekognition

- Custom image classification or object detection
- Easy to use with no ML expertise
- Fast experimentation

Want to build your own demo?

Code and instructions available on
GitHub. Build your own demo for
your local sign language!



Thank you!