AWS re:Invent

AIM360

Build a predictive maintenance system with Amazon SageMaker

Emily Webber
Machine Learning Specialist
Amazon Web Services

Laith Al-Saadoon
Senior Solutions Architect
Amazon Web Services

AWS re:Invent

© 2019, Amazon Web Services, Inc. or its affiliates. All rights reserved.

aws

Agenda

- Economics
- Architecture
- Workshop

Predictive maintenance economics

© 2019, Amazon Web Services, Inc. or its affiliates. All rights reserved.

What if I sell … motorcycles?

Selling motorcycles

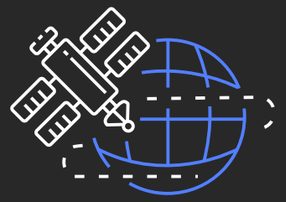

Forecast

Identify demand

Outreach

Contact buyers

Produce

Build the vehicles

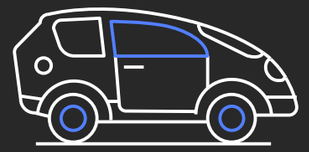

Deliver

Ship to customers

Monitor

Is it healthy?

Machine learning

can improve every step of this cycle

Selling motorcycles

Forecast

Identify demand

- Look at historical sales
- Combine economic levels
- Population projections
- Travel and traffic data
- Consider Amazon Forecast

Selling motorcycles

Outreach

Contact buyers

- Analyze sales data and cluster
- Combine browsing history
- Personal preferences
- Ranked search results
- Consider Amazon Personalize

Selling motorcycles

Produce

Build the vehicles

- Expensive machinery, which breaks!
- Manual labor: Train & supervise
- Optimize raw material supply
- Consider AWS IoT

Selling motorcycles

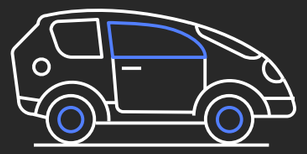

Deliver

Ship to customers

- Manual labor required to move
- Complex routing and scheduling
- Consider Amazon SageMaker

Selling motorcycles

Monitor

Is it healthy?

My customer’s motorcycle is almost certainly going to break

How do I predict that?

If I'm shipping 200,000 motorcycles every year,
and I can predict failure in even 10% of those,
and each detection saves me $200

200,000 X .1 x $200 = $4 million

I'm saving $4 million each year
from a single ML model and alerting system

Workshop setup

© 2019, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Setup instructions

https://bit.ly/2Dj37NF

1. Log into the Event Engine
2. Deploy the stack into your AWS account
3. Run predictions against your endpoint to get **RUL**
4. Use Ground Truth labelling to perform root cause analysis
5. Train a model to predict the type of failure

Architecture

© 2019, Amazon Web Services, Inc. or its affiliates. All rights reserved.

System design

- Say you have 100 engines
  - Some fail, some don’t
- You’re capturing all of the fields about those engines for each cycle
- You know exactly when the historical engines have failed

Let’s build a model to tell us the **remaining useful life** of each engine

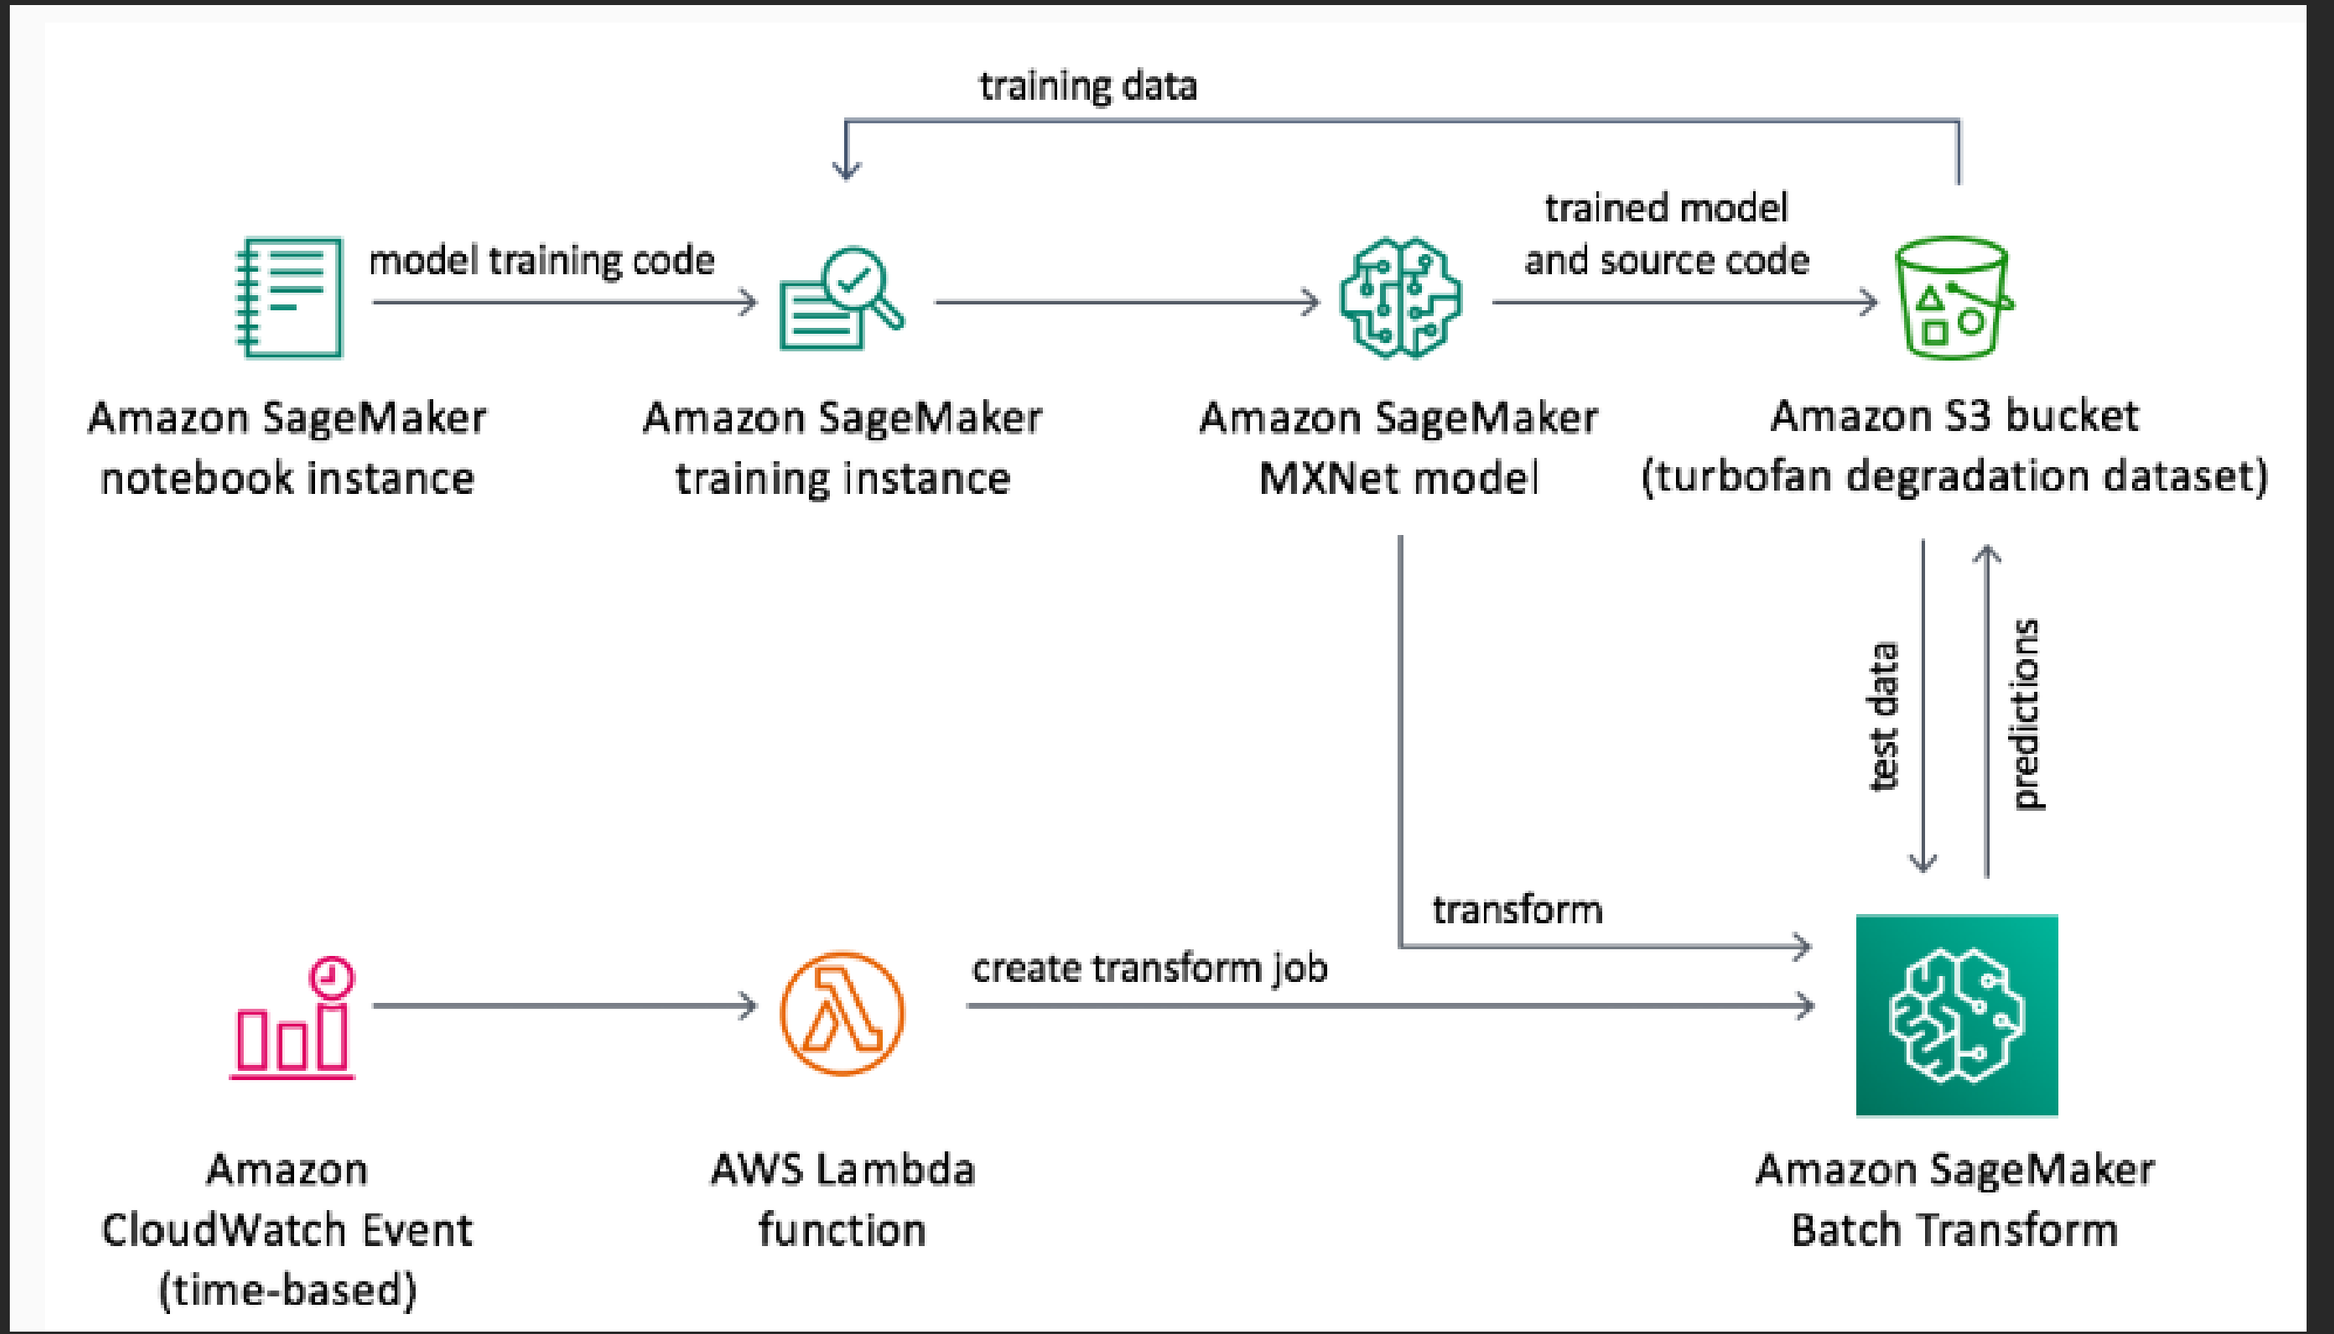
training data
model training code
trained model
and source code
Amazon SageMaker
notebook instance
Amazon SageMaker
training instance
Amazon SageMaker
MXNet model
Amazon S3 bucket
(turbofan degradation dataset)
test data
predictions
transform
create transform job
Amazon
CloudWatch Event
(time-based)
AWS Lambda
function
Amazon SageMaker
Batch Transform

Now, what if you could perform root cause analysis?

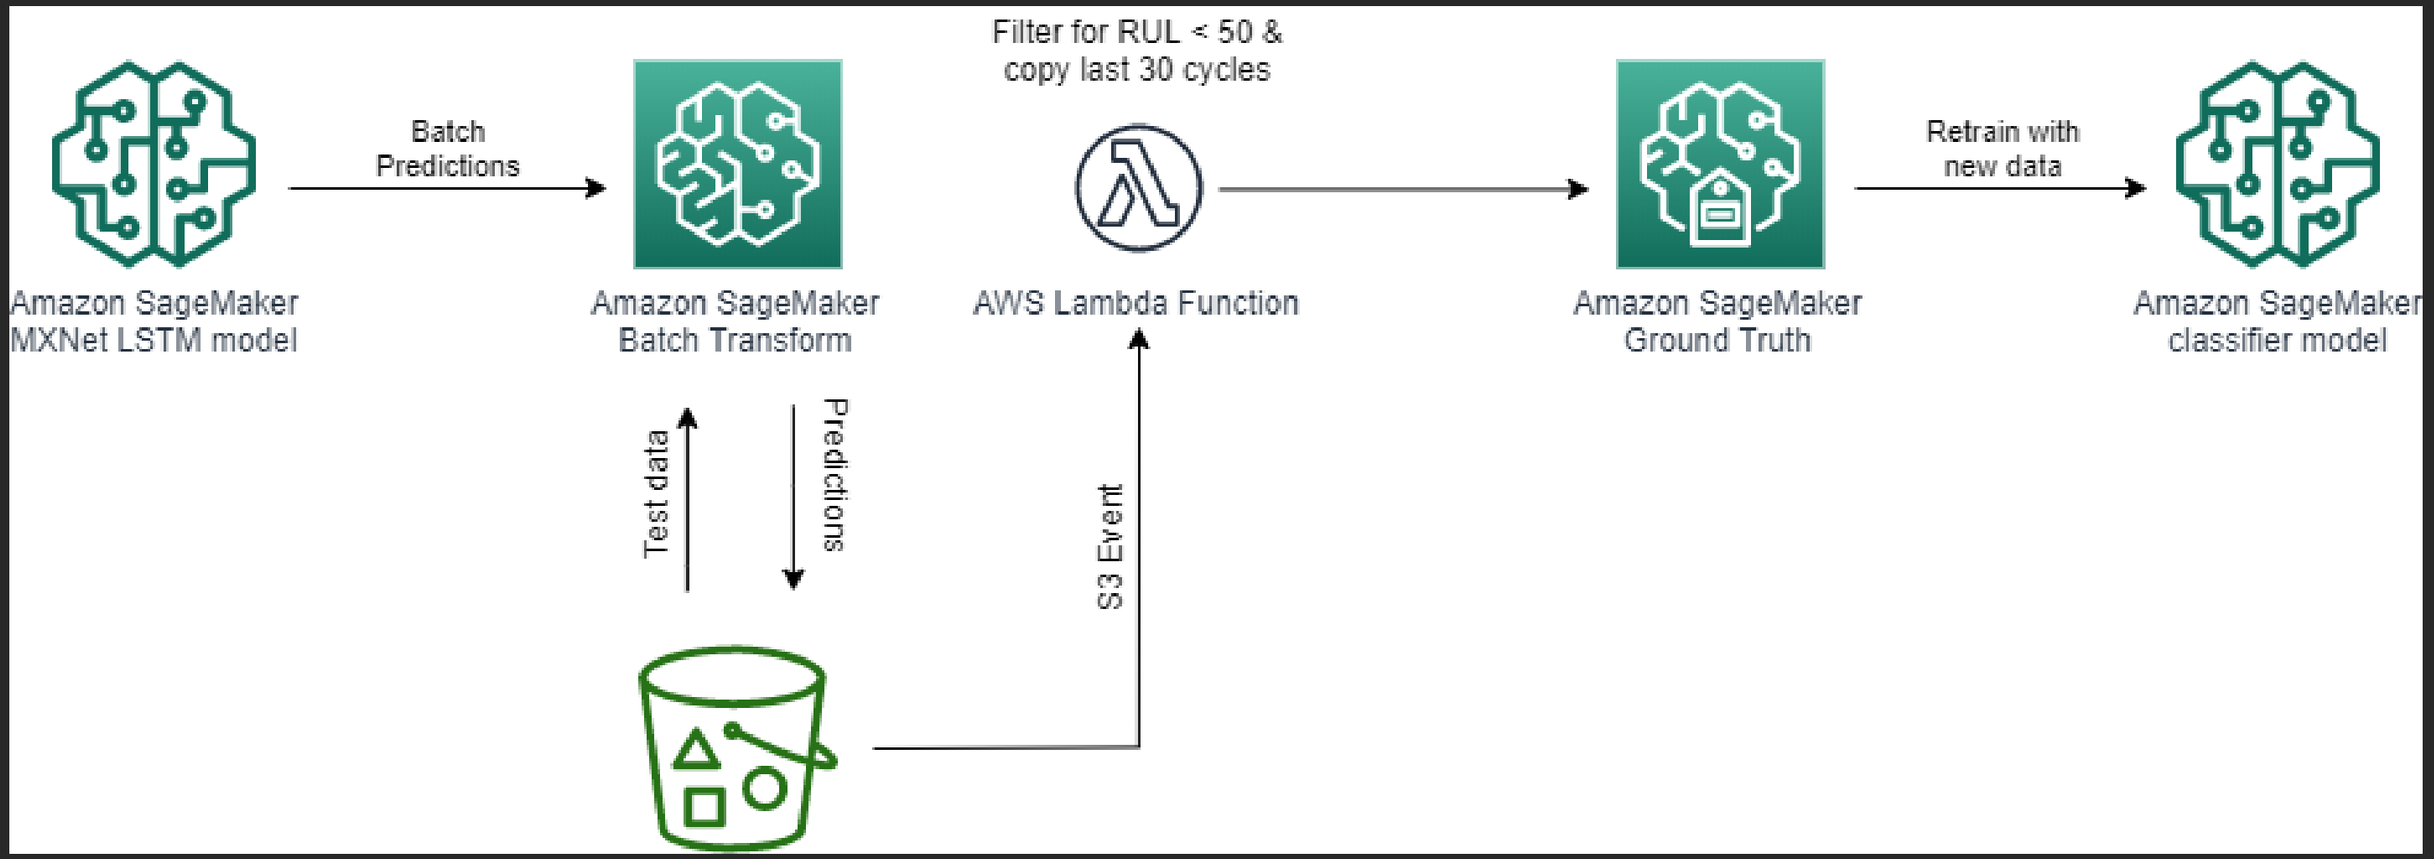

Learn ML with AWS Training and Certification

The same training that our own developers use, now available on demand

Role-based ML learning paths for developers, data scientists, data platform engineers, and business decision makers

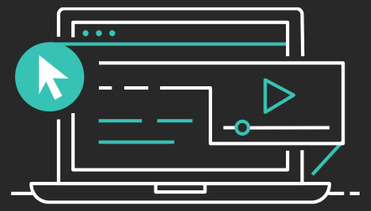

70+ free digital ML courses from AWS experts let you learn from real-world challenges tackled at AWS

Validate expertise with the
AWS Certified Machine Learning - Specialty exam

Visit https://aws.training/machinelearning

© 2019, Amazon Web Services, Inc. or its affiliates. All rights reserved.

aws training and certification

Thank you!

Emily Webber

https://www.linkedin.com/in/emily-webber-921b4969/

Laith Al-Saadoon

https://www.linkedin.com/in/laithalsaadoon/

AWS re:Invent

© 2019, Amazon Web Services, Inc. or its affiliates. All rights reserved.

aws

Please complete the session survey in the mobile app.

© 2019, Amazon Web Services, Inc. or its affiliates. All rights reserved.