# Real-Time Decisions Using ML on the Google Cloud Platform

**Przemysław Pastuszka & Carlos Garcia**
QCon London
7th March 2018

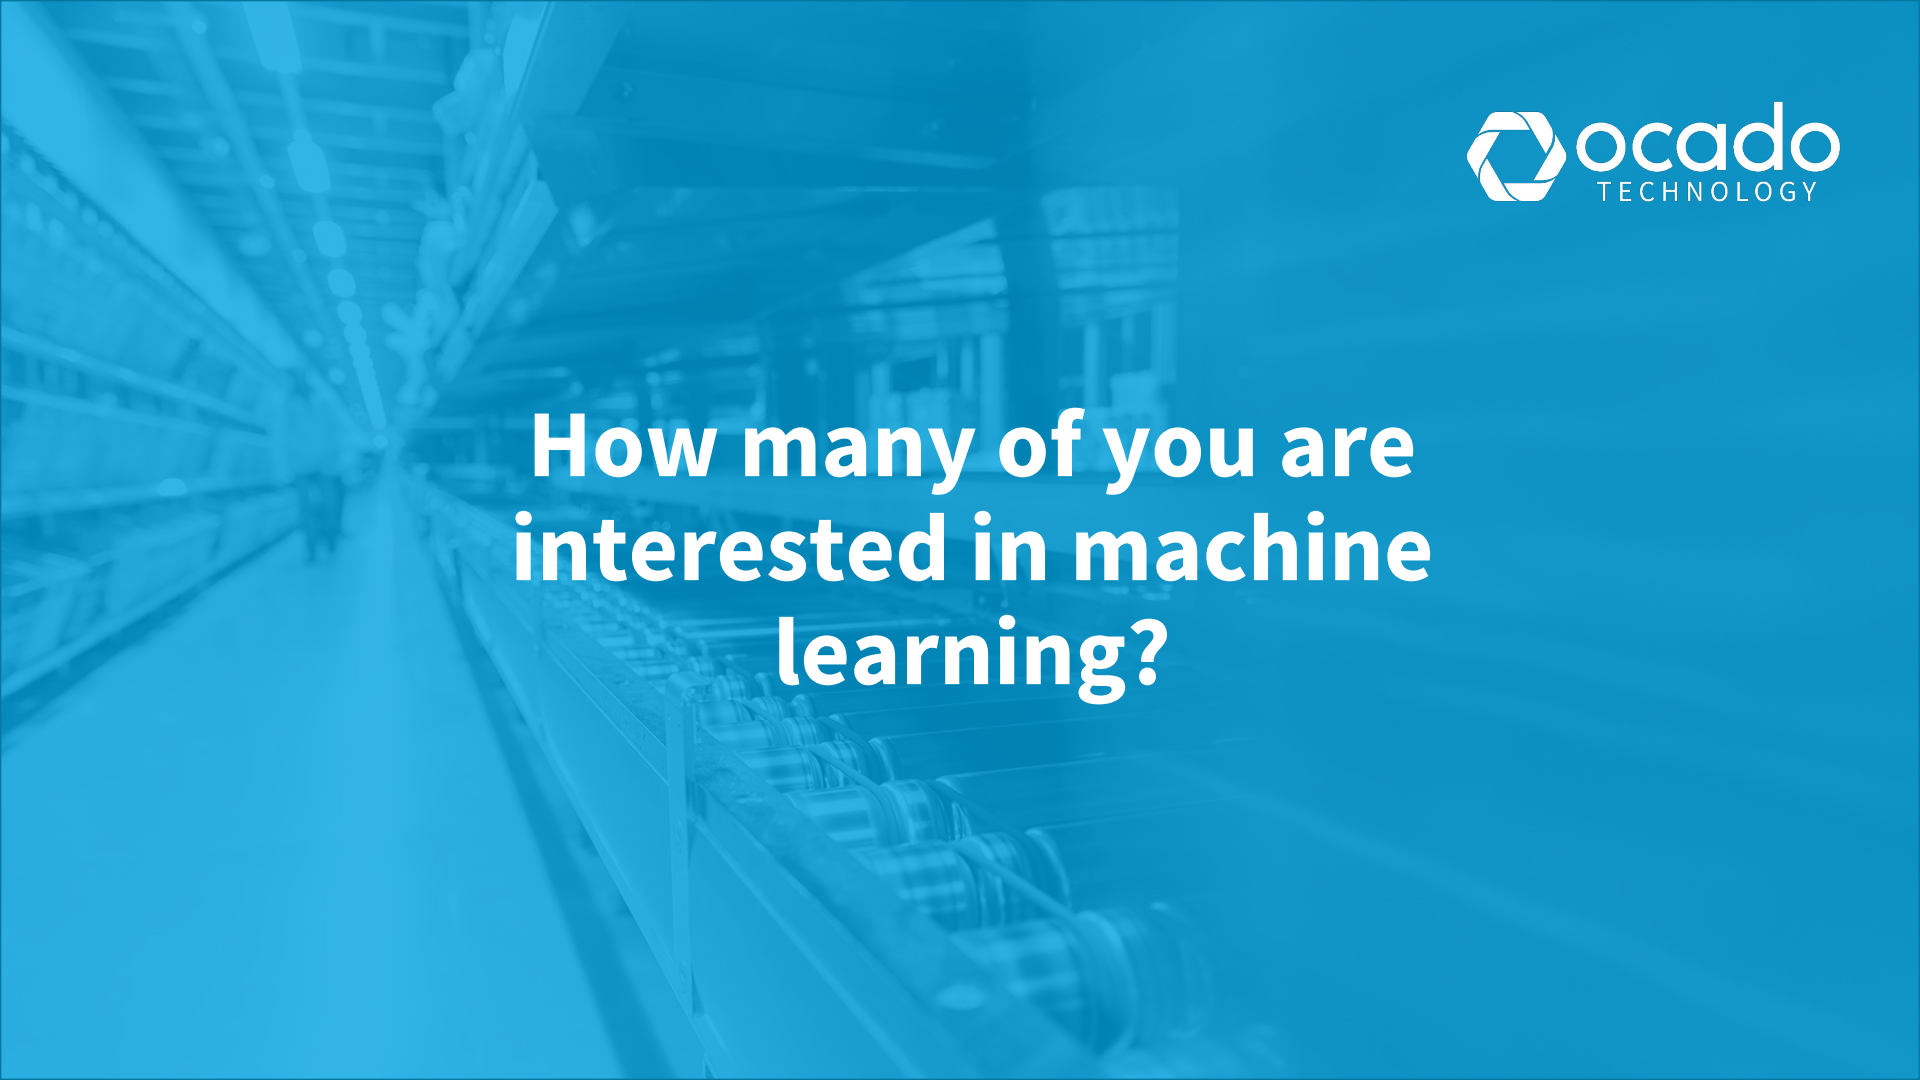


# How many of you are interested in machine learning?

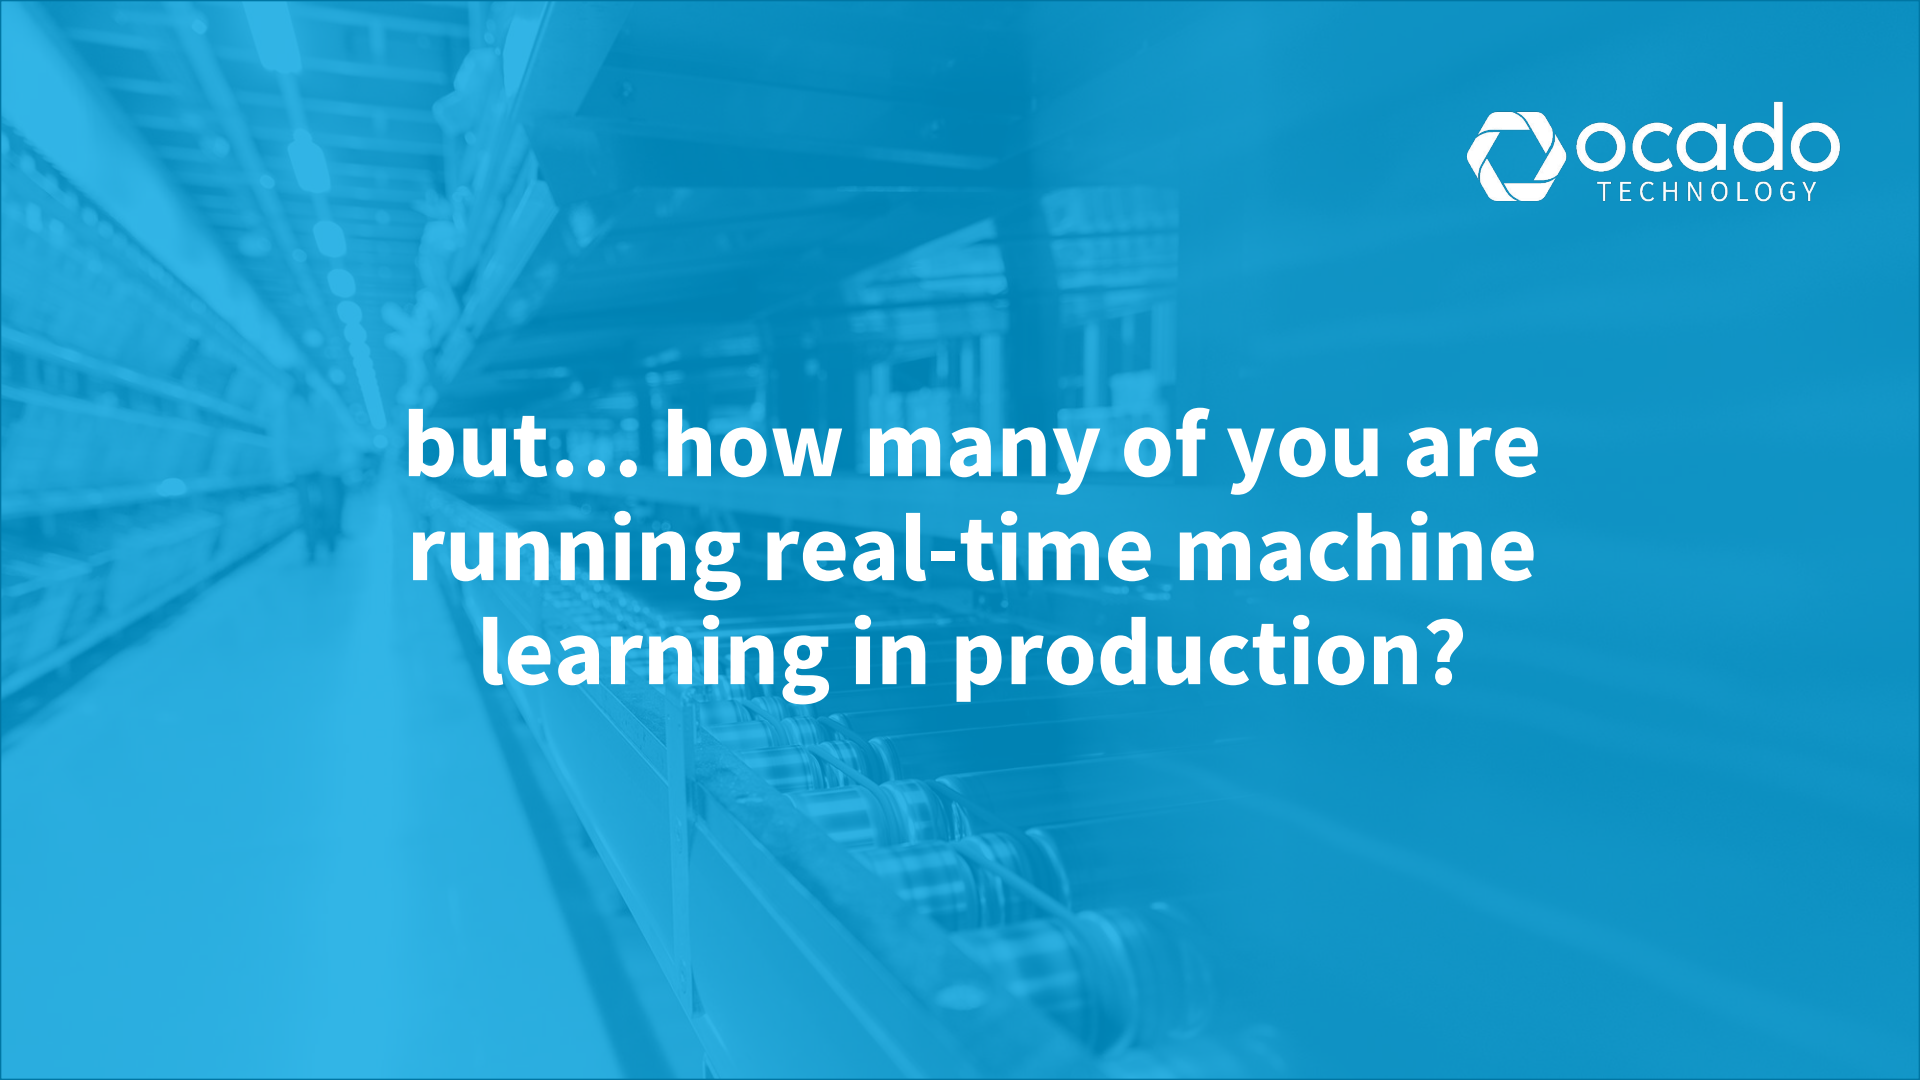
ocado
TECHNOLOGY
but… how many of you are running real-time machine learning in production?

# Who is Ocado?

ocado TECHNOLOGY

Ocado is the **world's largest** dedicated online grocery retailer

We have **645,000** active shoppers

And **49,000** SKUs in our webshop

Three highly-automated fulfilment centres

**263,000** orders a week 'picked'

**3 million** routing calculations per second

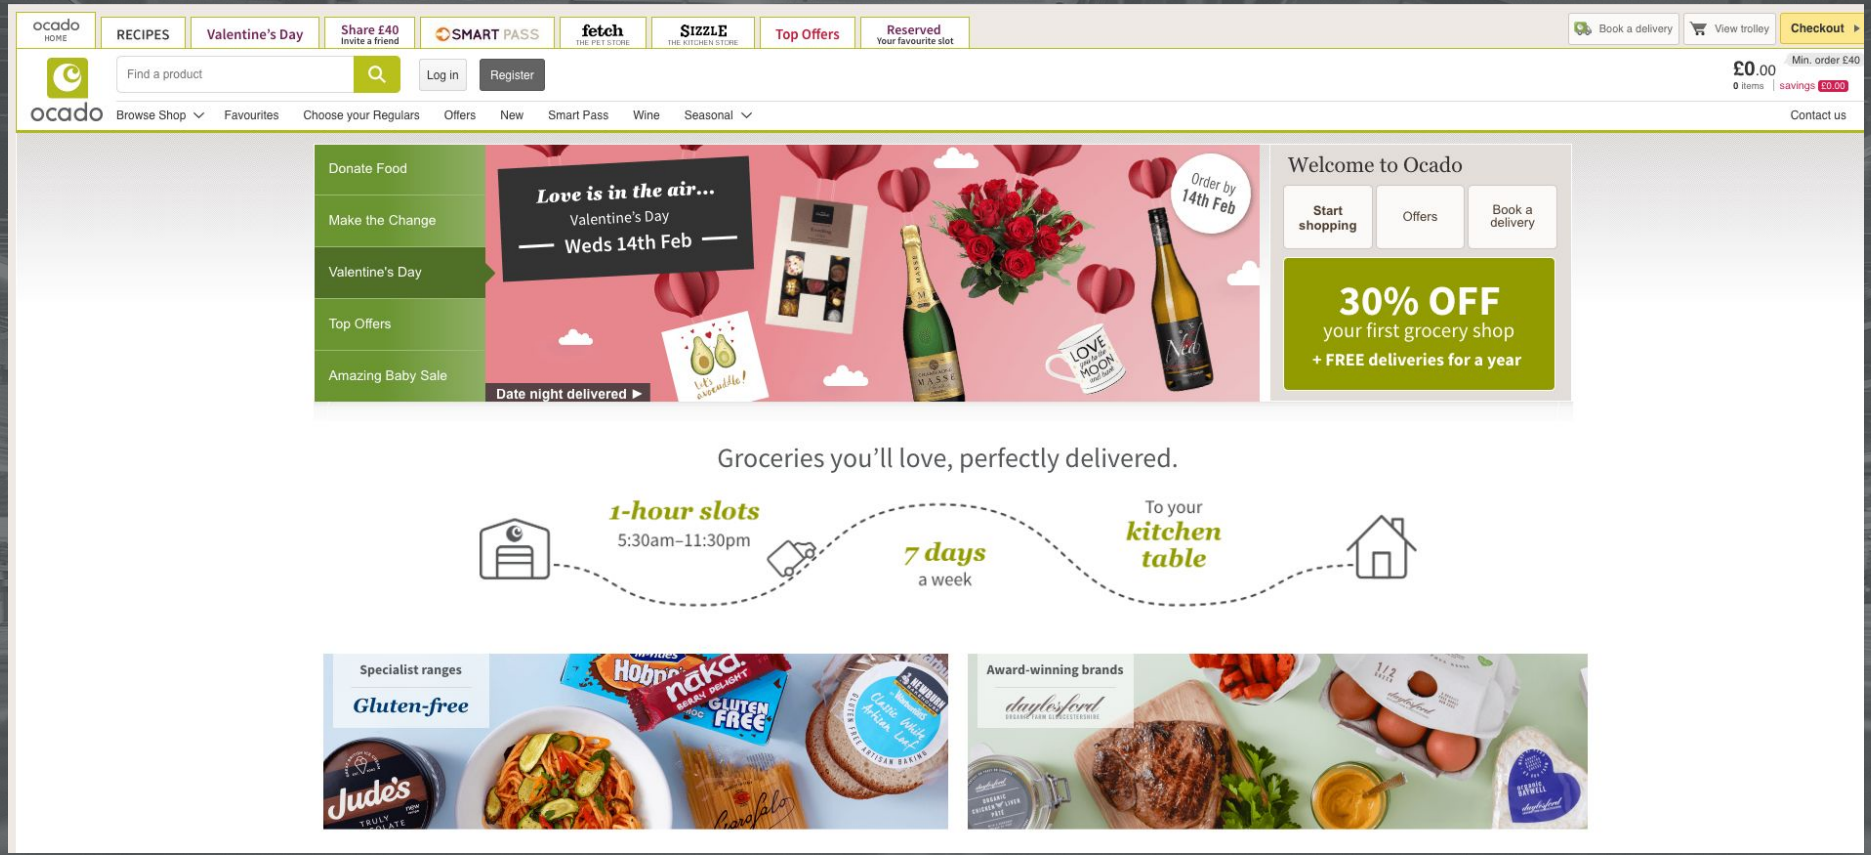
ocado HOME
RECIPES
Valentine's Day
Share £40 Invite a friend
SMART PASS
fetch THE PET STORE
SIZZLE THE KITCHEN STORE
Top Offers
Reserved Your favourite slot
Book a delivery
View trolley
Checkout
Find a product
Log in
Register
£0.00 0 items
Min. order £40
savings £0.00
ocado
Browse Shop
Favourites
Choose your Regulars
Offers
New
Smart Pass
Wine
Seasonal
Contact us
Donate Food
Make the Change
Valentine's Day
Top Offers
Amazing Baby Sale
Love is in the air...
Valentine's Day
Weds 14th Feb
Order by 14th Feb
Date night delivered ▶
Welcome to Ocado
Start shopping
Offers
Book a delivery
30% OFF
your first grocery shop
+ FREE deliveries for a year
Groceries you'll love, perfectly delivered.
1-hour slots
5:30am–11:30pm
7 days
a week
To your
kitchen
table
Specialist ranges
Gluten-free
Award-winning brands
daylesford

# What Ocado Technology does

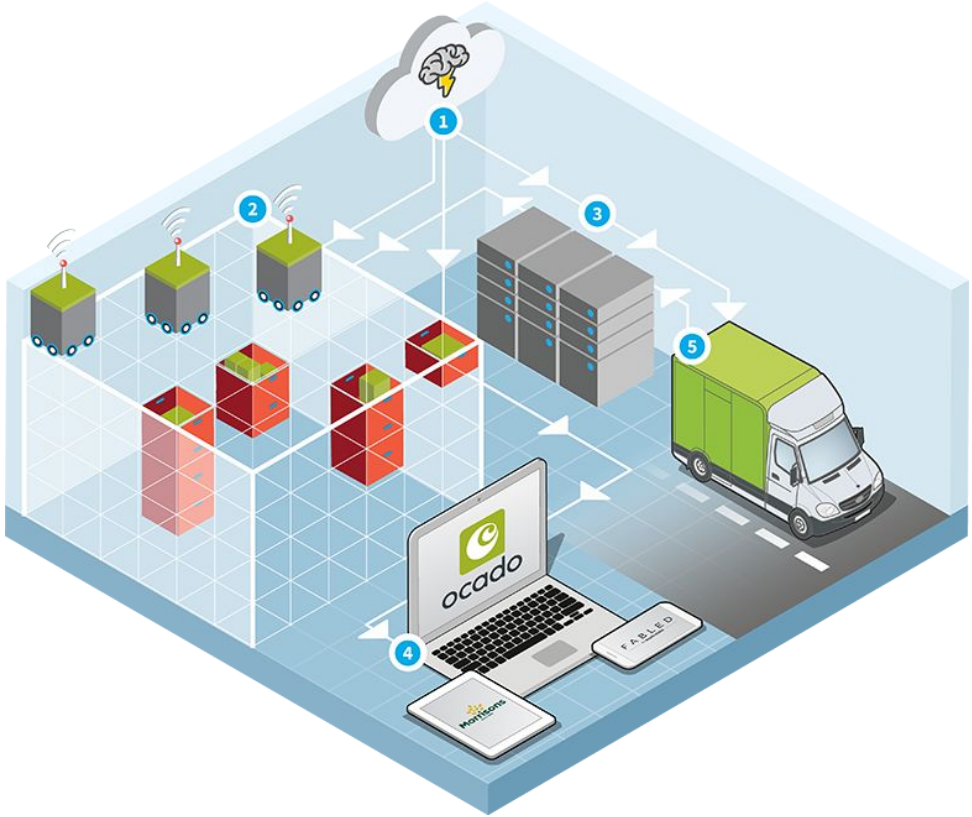


(1) **Cloud and AI**
(2) **Automation and robotics**
(3) **Big Data**
(4) **Web and app development**
(5) **IoT**

# Fraud: An ML journey

ocado TECHNOLOGY

ocado TECHNOLOGY


ocado TECHNOLOGY

# But then… what is fraud?

- Mainly chargebacks
- Other types of fraud?
- Learn from the actual outcome

# Do I really need to do any ML?

# Know your target

- Do you need ML?
- What do you want to predict?
- How good are you at predicting that?

# Cost of mistakes

- False positives and false negatives
- How expensive are they?

# Start with heuristics

- Ask domain experts
- Derive rules from expert knowledge
  - “If more than 80% of order is alcohol, then classify as risky”

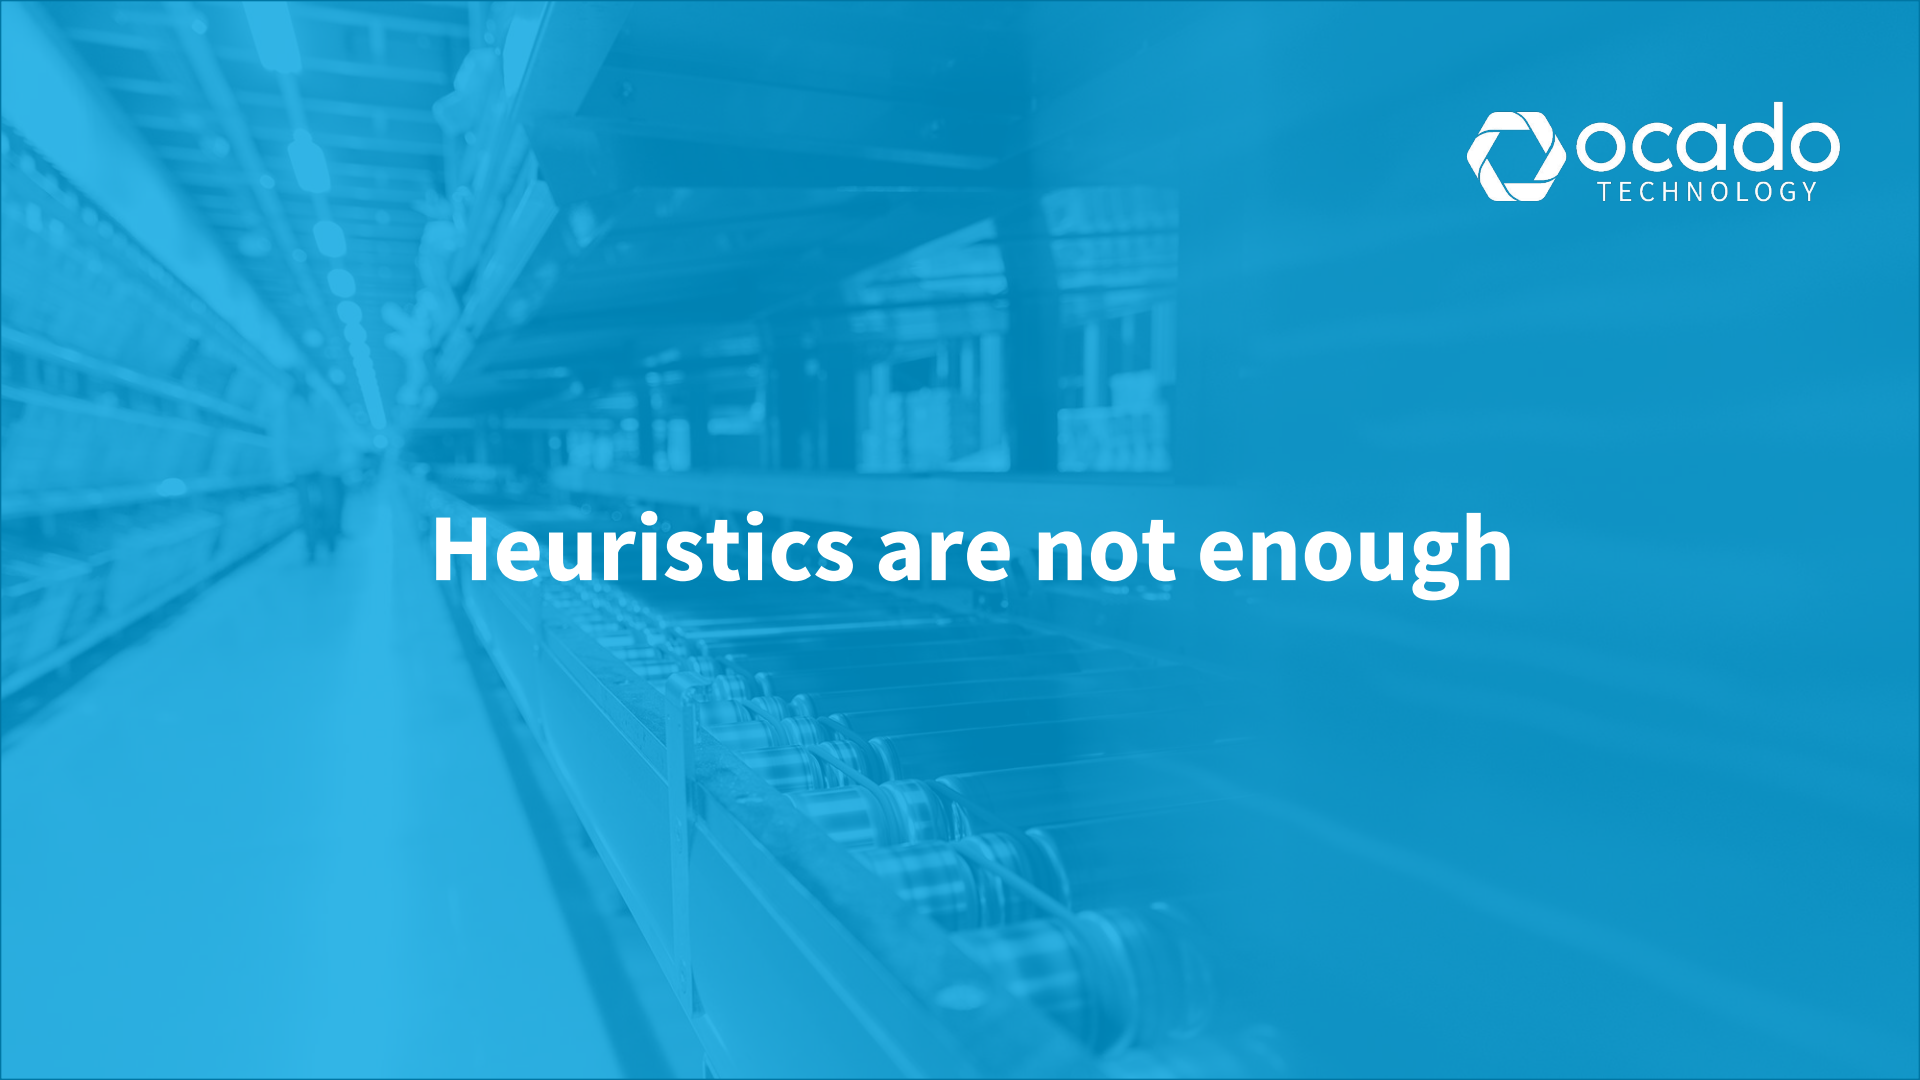

# Heuristics are not enough

# Motivations for Machine Learning

ocado TECHNOLOGY

Data-driven

# Motivations for Machine Learning

ocado TECHNOLOGY

Data-driven

Fraudsters learn

# Motivations for Machine Learning

ocado TECHNOLOGY

Data-driven

Fraudsters learn

Customer patterns

# Motivations for Machine Learning

ocado TECHNOLOGY

Data-driven

Fraudsters learn

Customer patterns

Business changes

# Challenges

- Fraud (human) agents
- ML is affected by human decisions
- Unbalanced classes (fraud / not-fraud)


# What ML model do you choose?

ocado TECHNOLOGY

**“With great power there must also come… great responsibility”**

Spider-Man

# Criteria

Online vs batch predictions

# Criteria

Online vs batch predictions

Explainable predictions

# Challenge your explanations

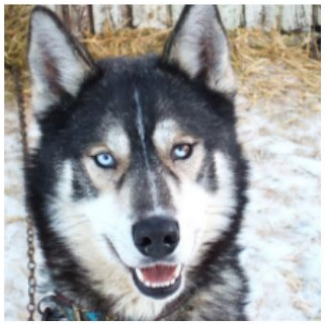

(a) Husky classified as wolf

(b) Explanation

"Why should I trust you?"

2016, M. Tulio, S. Singh, C. Guestrin

ocado TECHNOLOGY

# Criteria

Online vs batch predictions

Explainable predictions

Programming language

# Criteria

Online vs batch predictions

Explainable predictions

Programming language

Cloud vs on-premise

# Our ML choice

ocado TECHNOLOGY

# Criteria

ocado TECHNOLOGY

| | |
|---|---|
| Online vs batch predictions | online |
| Explainable predictions | preferable |
| Programming language | Python / Java |
| Cloud vs on-premise | cloud |

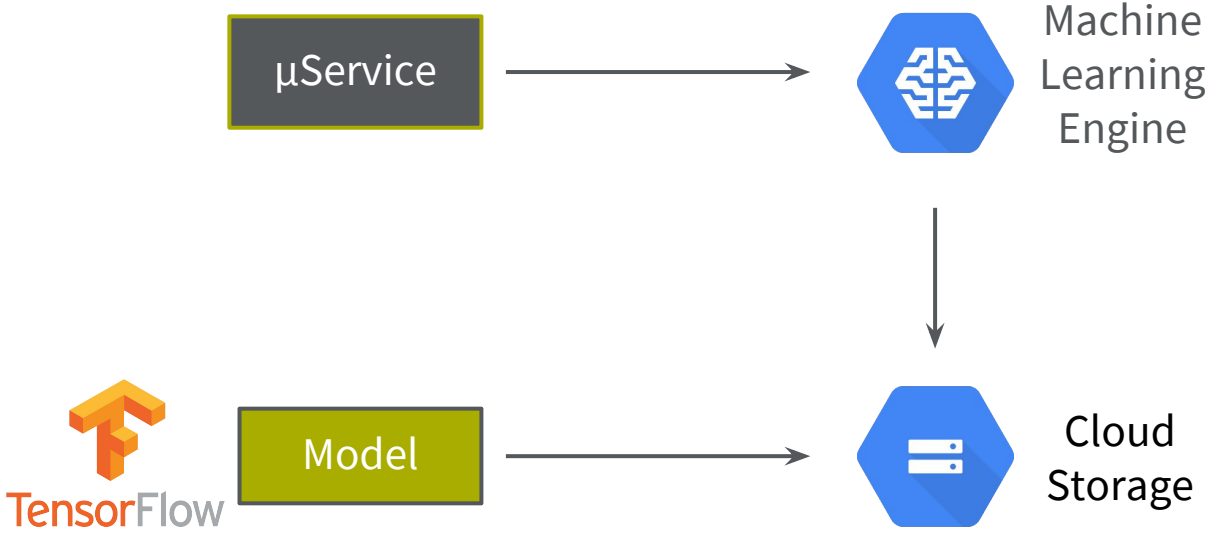


ocado TECHNOLOGY

TensorBoard SCALARS GRAPHS PR CURVES INACTIVE

Time Display Type

STEP RELATIVE WALL

colors
step 9

mask_every_other_prediction
step 9

Runs

Write a regex to filter runs

- colors
- mask_every_other_prediction
- temperature:t0=270,tA=270,kH=0.001
- temperature:t0=270,tA=270,kH=0.005
- temperature:t0=270,tA=310,kH=0.001
- temperature:t0=270,tA=310,kH=0.005
- temperature:t0=270,tA=350,kH=0.001
- temperature:t0=270,tA=350,kH=0.005
- temperature:t0=310,tA=270,kH=0.001
- temperature:t0=310,tA=270,kH=0.005
- temperature:t0=310,tA=310,kH=0.001
- temperature:t0=310,tA=310,kH=0.005
- temperature:t0=310,tA=350,kH=0.001
- temperature:t0=310,tA=350,kH=0.005
- temperature:t0=350,tA=270,kH=0.001

TOGGLE ALL RUNS

/Users/chizeng/Desktop/pr_curve_demo

.*

Tags matching /.*/ (all tags) 3

classifying blue
tag: blue/pr_curves

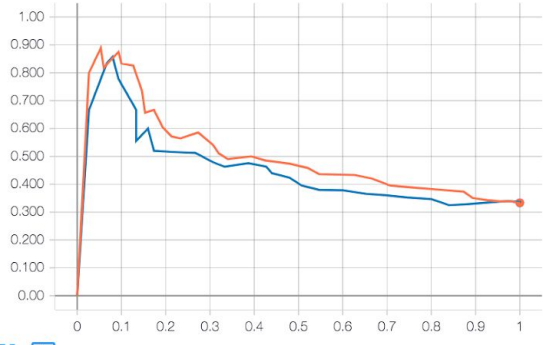


colors is at **step 9**
(Sat Aug 26 2017 00:34:34 GMT-0400 (EDT))
mask_every_other_prediction is at **step 9**
(Sat Aug 26 2017 00:34:35 GMT-0400 (EDT))

classifying green
tag: green/pr_curves

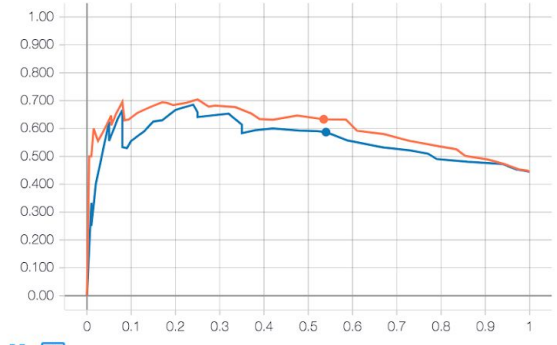


| Run | Threshold | Precision | Recall | TP | FP | TN | FN |
|---|---|---|---|---|---|---|---|
| colors | 0.3000 | 0.6331 | 0.5350 | 107 | 62 | 188 | 93 |
| mask_every_other_prediction | 0.2600 | 0.5870 | 0.5400 | 54 | 38 | 87 | 46 |

classifying red
tag: red/pr_curves

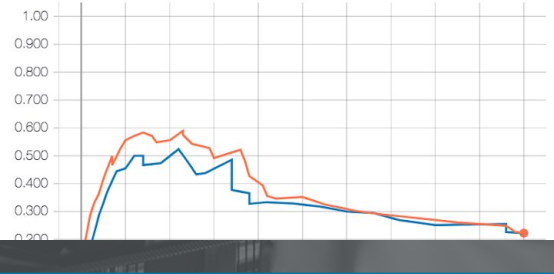

# Interesting alternatives

Amazon
SageMaker





# Interesting alternatives

Amazon SageMaker





Google Cloud Machine Learning Engine





# Problem #1
# Not fast enough

# Data exploration cycle

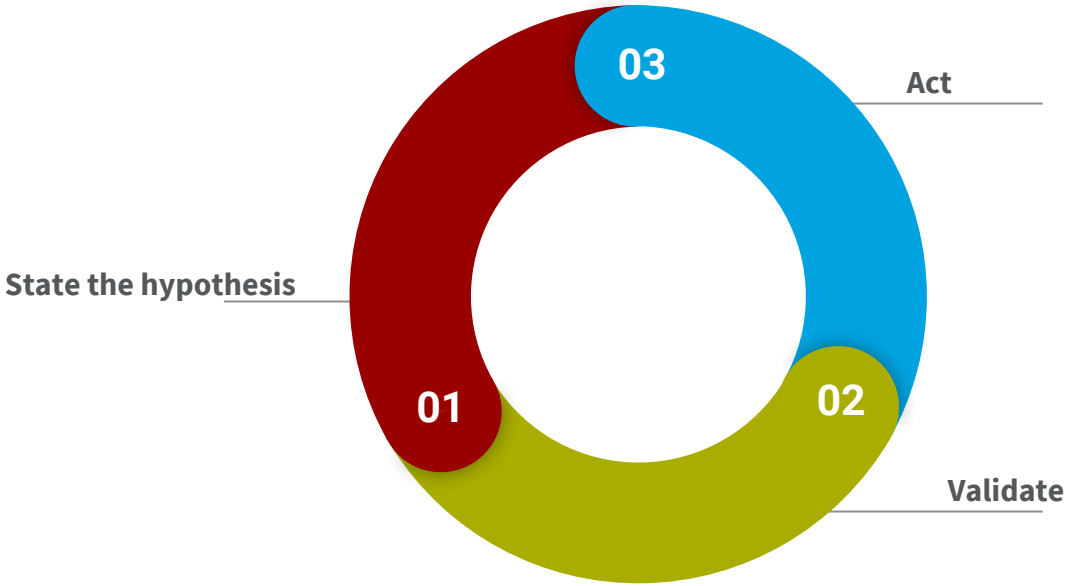


# Validate your hypothesis - fast!

Big Query

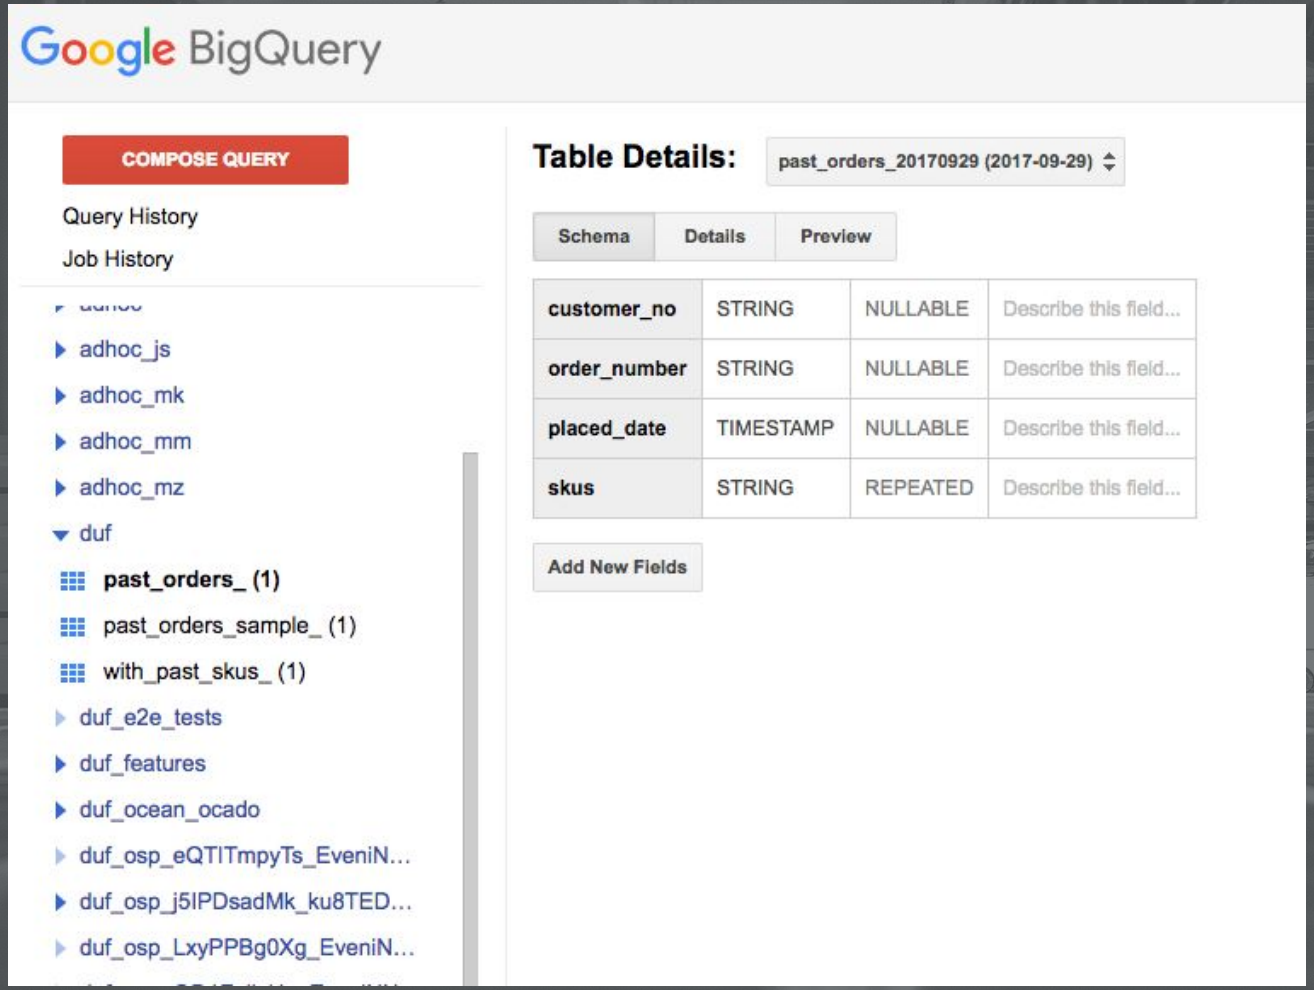


ocado TECHNOLOGY

New Query

```
1  SELECT * FROM [rsh-osp-eu-mlservices:duf.past_orders_20170929]
```

RUN QUERY | Save Query | Save View | Format Query | Show Options

**Table Details:** past_orders_20170929 (2017-09-29)

Schema | Details | Preview

| Row | customer_no | order_number | placed_date | skus |
|---|---|---|---|---|
| 1 | 00627486 | 3838642622 | 2017-09-12 00:00:00.000 UTC | 330215011 |
| | | | | 45547011 |

ocado TECHNOLOGY


# Problem #2
# Data delivered too late

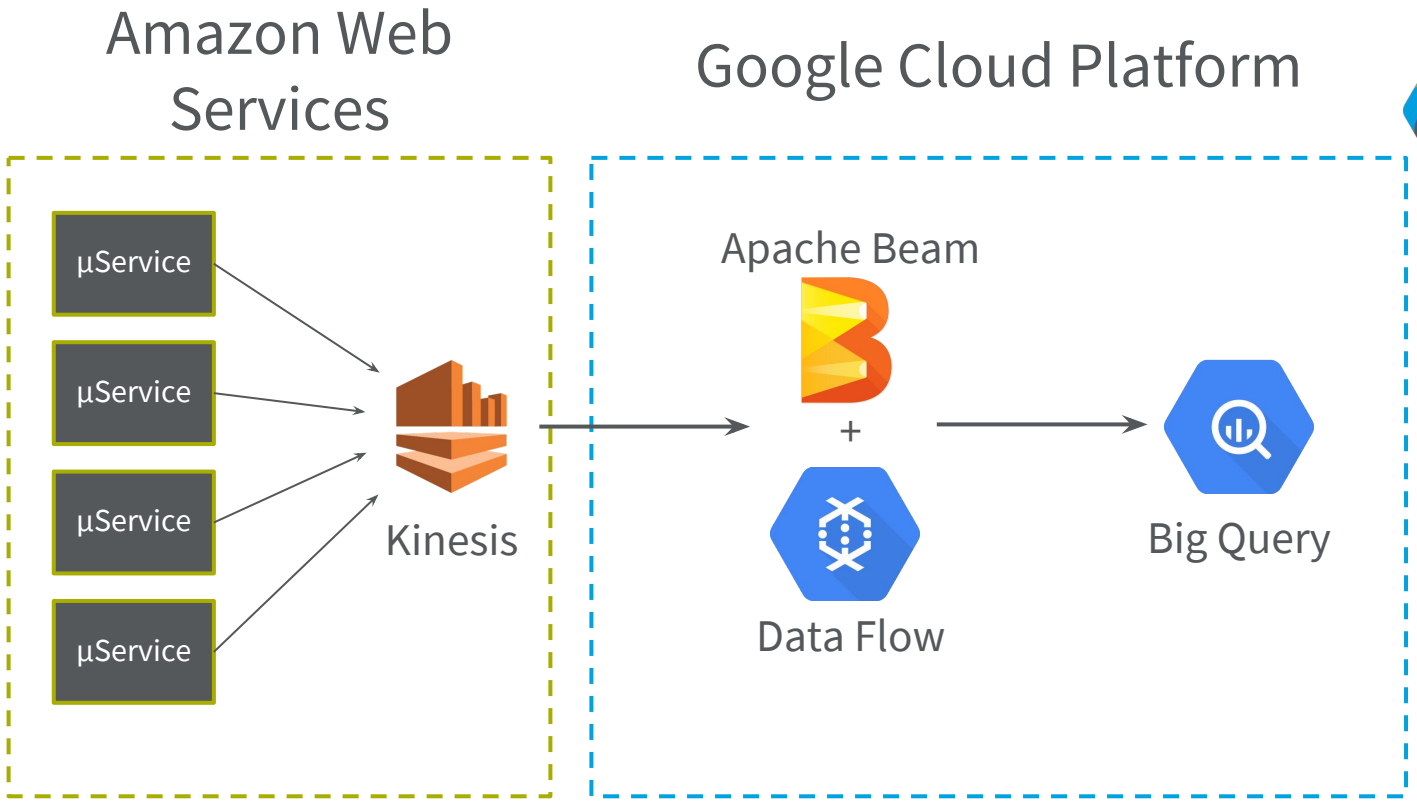


ocado TECHNOLOGY

```
List<String> strings = ...
strings.stream().collect(
    Collectors.groupingBy(
        word -> word.charAt(0),
        Collectors.counting()));
```

```
PCollection<String> pipeline = ...
pipeline
    .apply(MapElements.via(row -> KV.of(word.charAt(0), word)))
    .apply(GroupByKey.create())
    .apply(Count.perKey())
```

ocado TECHNOLOGY

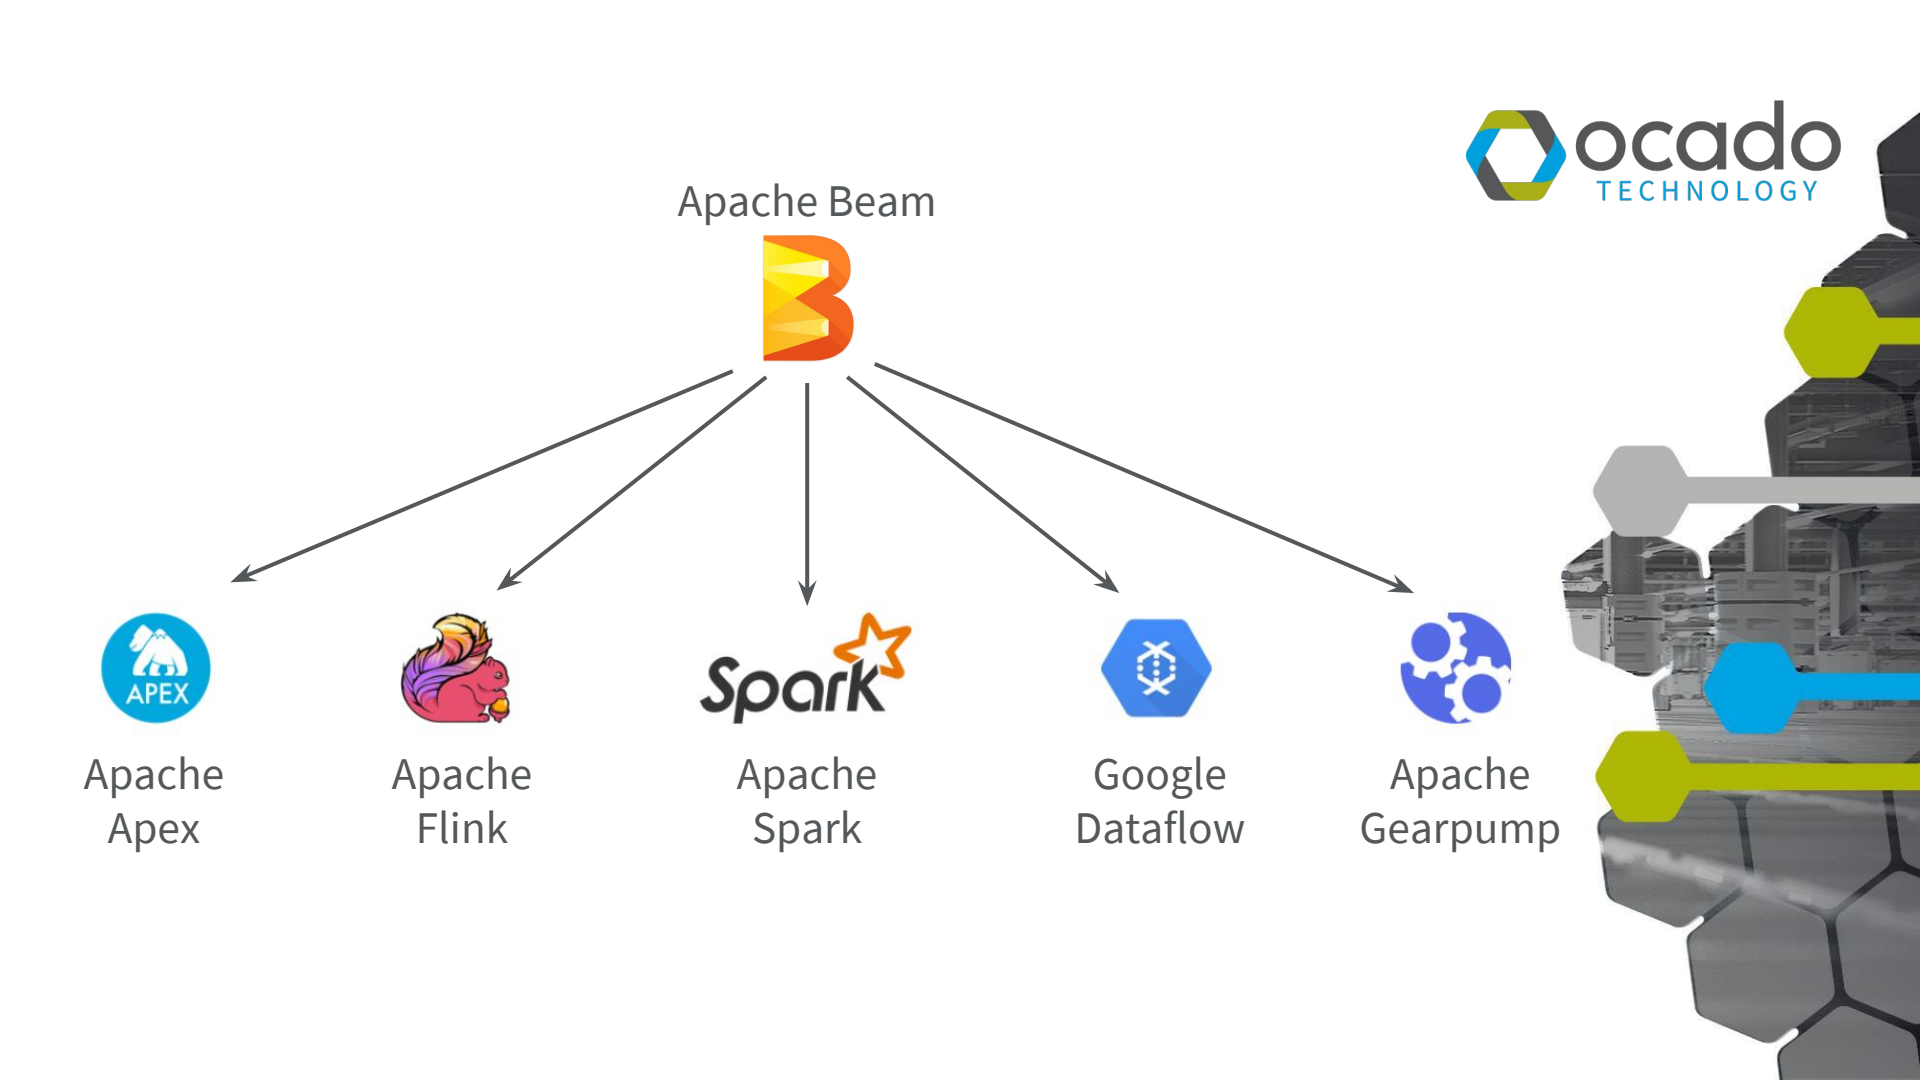
Apache Beam
APEX
Spark
Apache Apex
Apache Flink
Apache Spark
Google Dataflow
Apache Gearpump
ocado
TECHNOLOGY

# Problem #3
# Missing data

# Missing data



# Capture every change to the business state

# Training

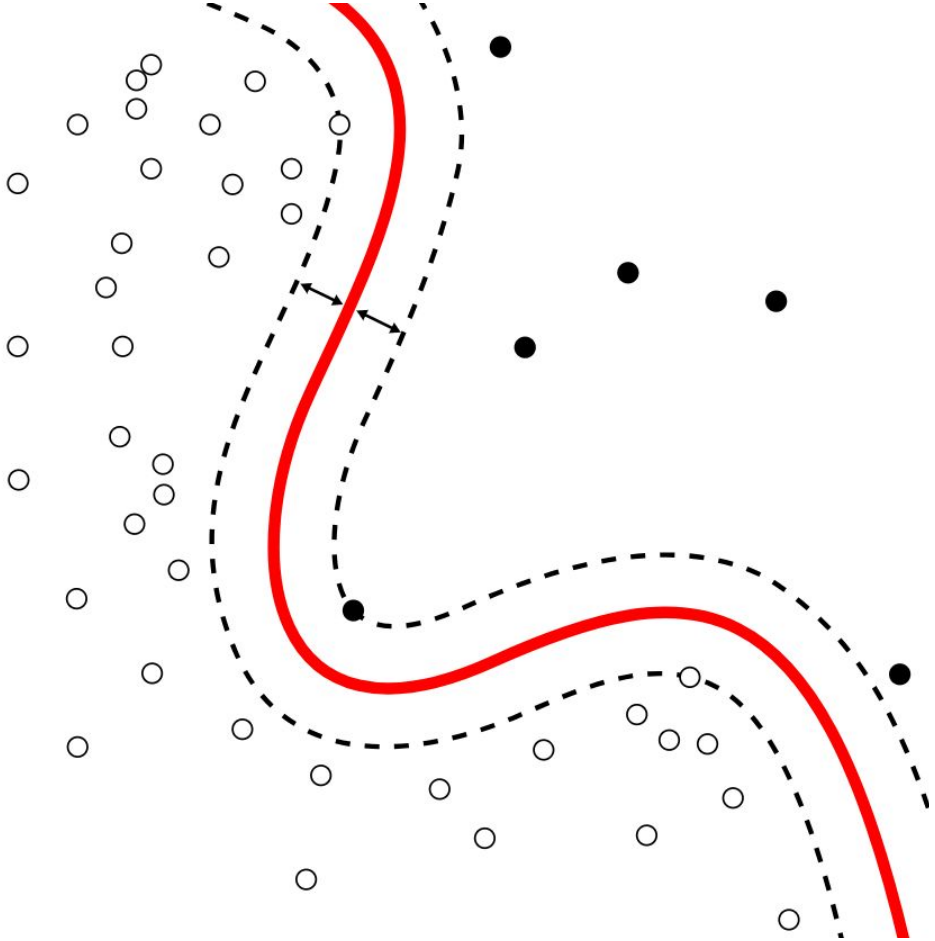

ocado TECHNOLOGY

train($C_1$, … $C_N$, $O_1$, … $O_N$, Y) = model

$C_1$, … $C_N$, $O_1$, … $O_N$ - customer and order features

$C_1$ - Average basket size for the customer

$O_1$ - % of alcoholic items in current order

…

Y - Fraud or not fraud

ocado TECHNOLOGY

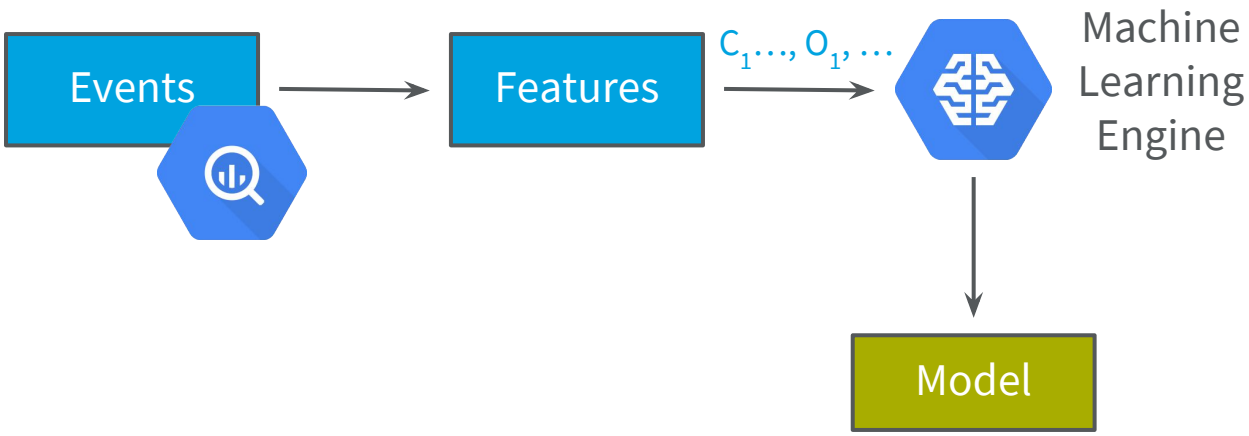


ocado TECHNOLOGY

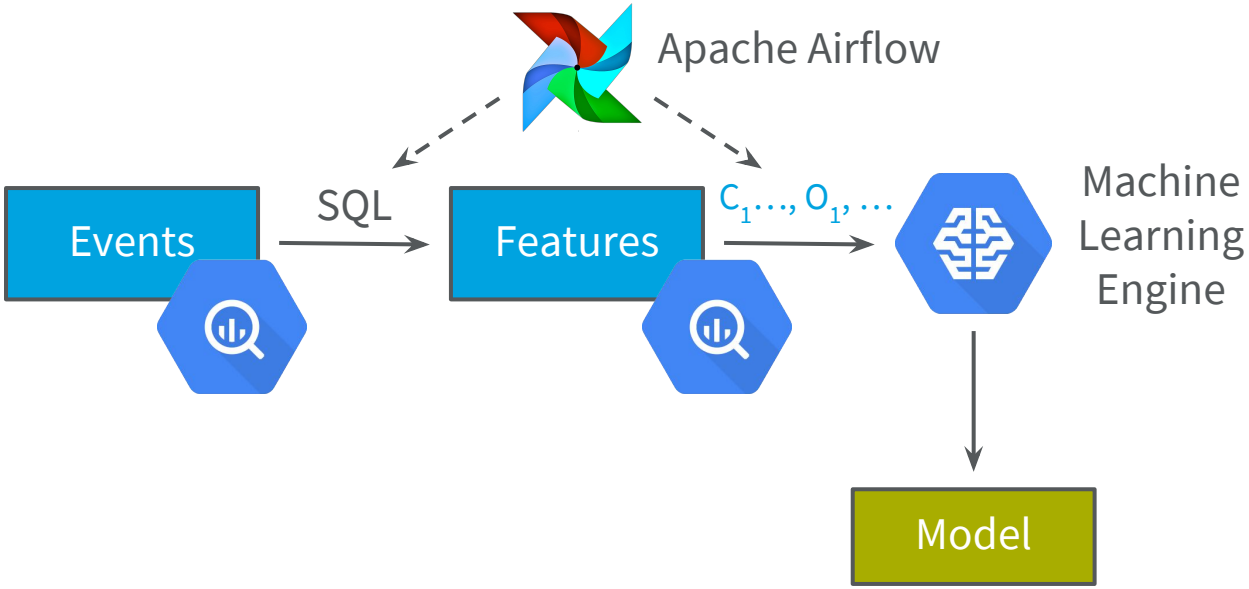


ocado TECHNOLOGY


# Serving predictions

train($C_1$, … $C_N$, $O_1$, … $O_N$, $Y$) = model

model($C_1$, … $C_N$, $O_1$, … $O_N$) = prediction

$C_1$, … $C_N$, $O_1$, … $O_N$ - customer and order features

$C_1$ - Average basket size for the customer

$O_1$ - % of alcoholic items in current order

…

$Y$ - Fraud or not fraud

prediction - Probability of current order being fraudulent

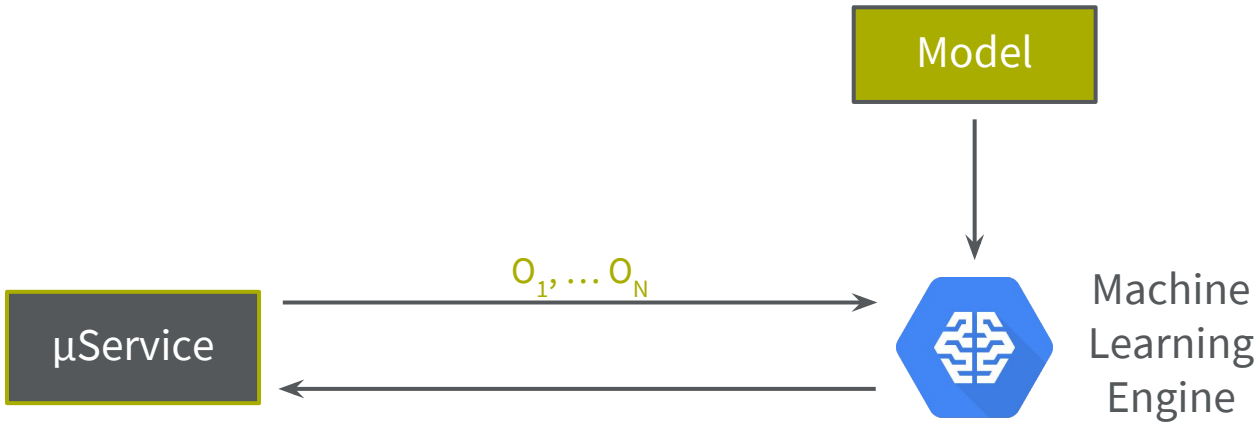
Model
$O_1, \ldots O_N$
µService
Machine
Learning
Engine

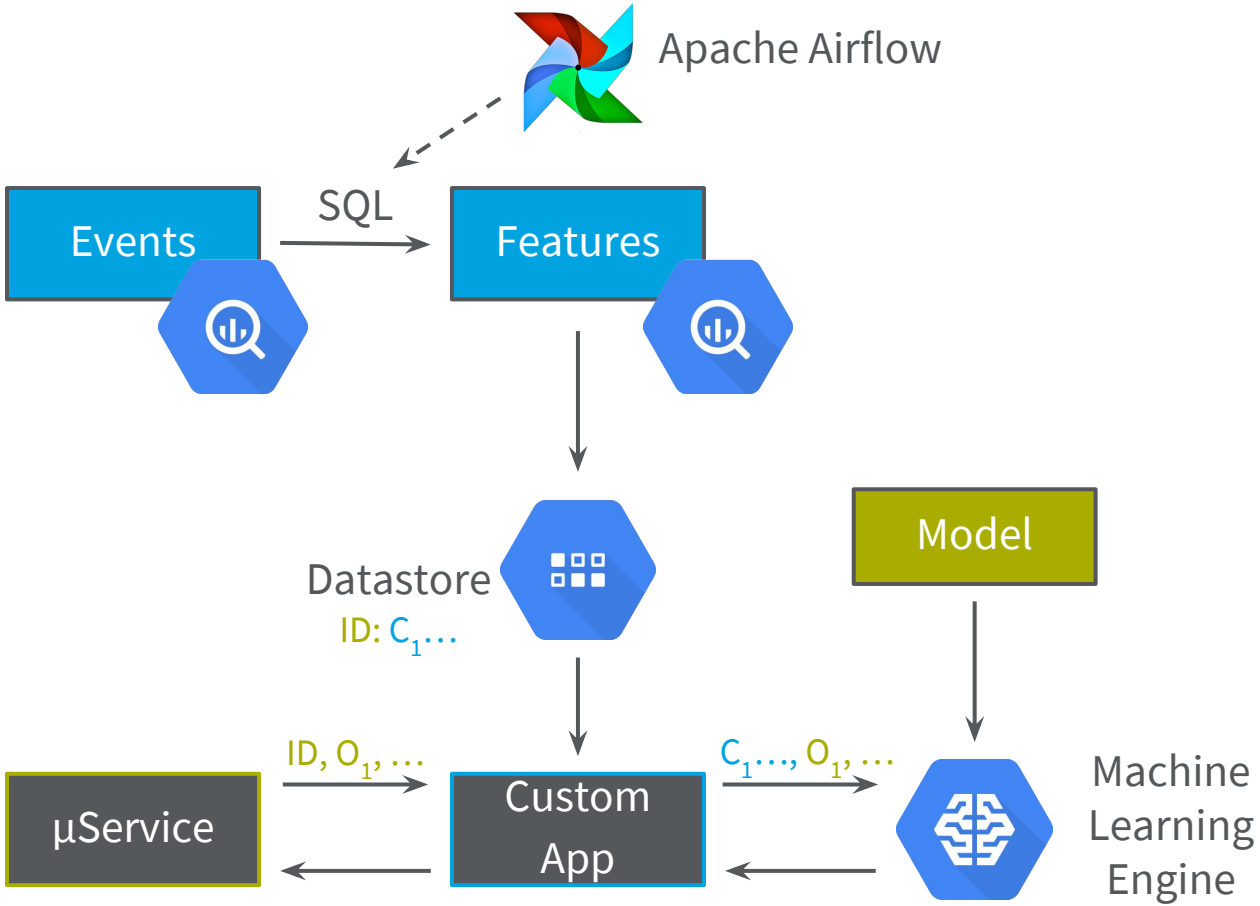


ocado TECHNOLOGY

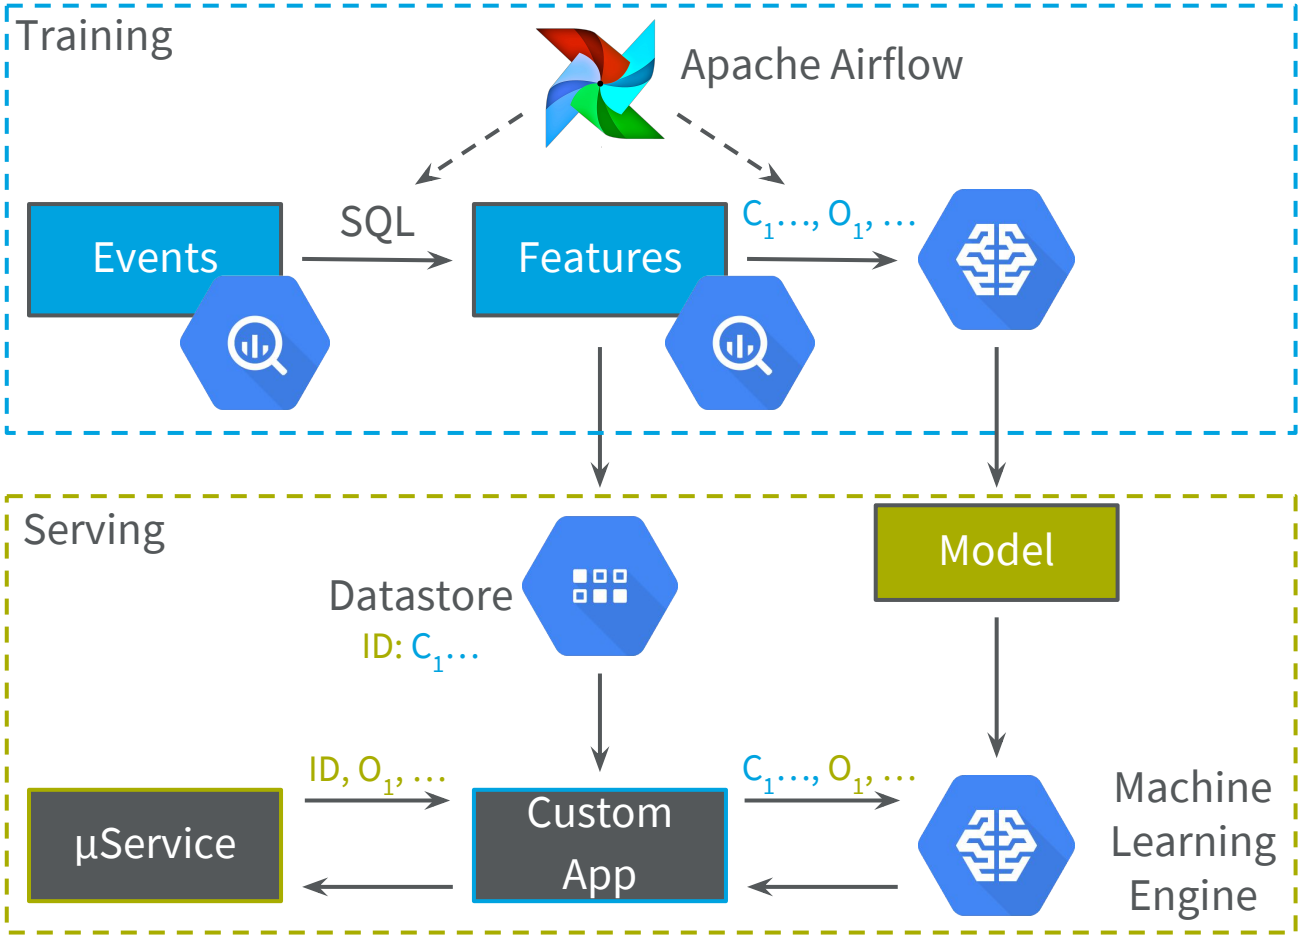


ocado TECHNOLOGY


# Architecting for the future

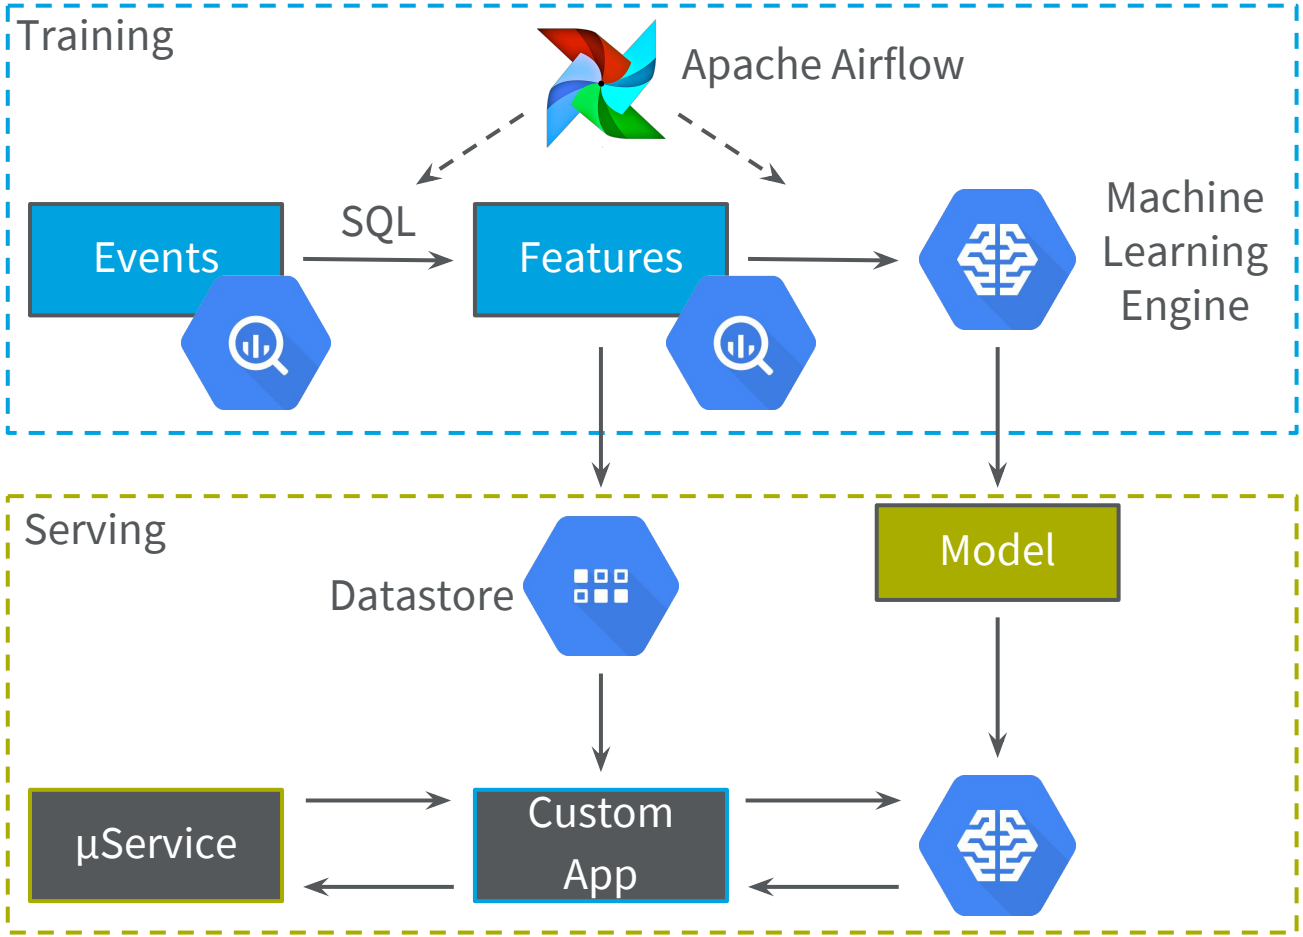


ocado TECHNOLOGY

Know your target

Keep It Simple

Choose your model wisely

Google Cloud ML Engine for Neural Nets

Have data and tools ready

BigQuery is king

Unified architecture for training and serving predictions

ocado TECHNOLOGY


ocado
TECHNOLOGY
Thank you!