zalando

Continuous Live Monitoring of Machine Learning Models with Delayed Label Feedback

Patrick Baier

Lorand Dali

Zalando Payments

June 2018

OUTLINE

Who we are and what we do

Why we should monitor

Prediction Monitoring

Our implementation

WHO WE ARE
AND
WHAT WE DO

WHO WE ARE

Patrick Baier

- Data Scientist at Zalando (~ 3.5 years)
- PhD in Computer Science from Uni Stuttgart

Lorand Dali

- Data Scientist at Zalando (~ 1.5 years)
- Diploma in Computer Science from the Technical University of Cluj Napoca

WHAT WE DO

Detect and prevent payment fraud

Shopping

Shipping

Payment

WHAT WE DO

Detect and prevent payment fraud

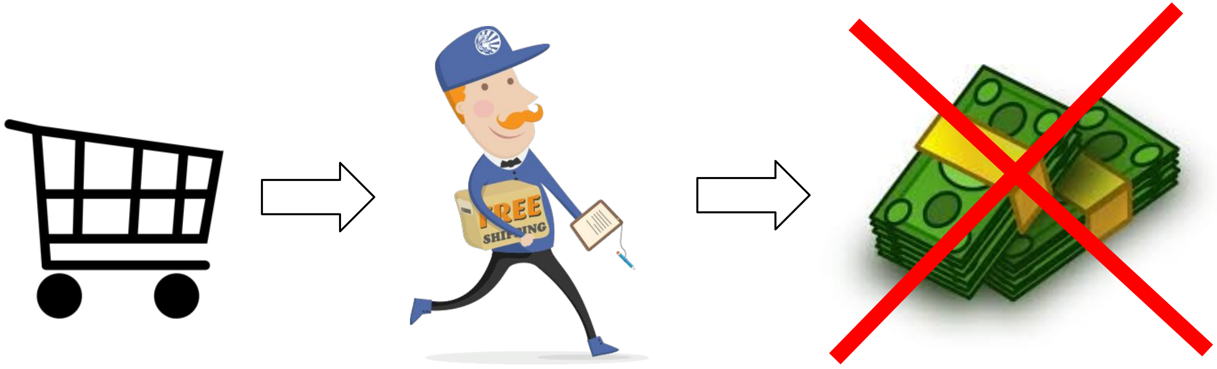

Shopping Shipping Payment

MODEL TRAINING

RUNTIME SYSTEM

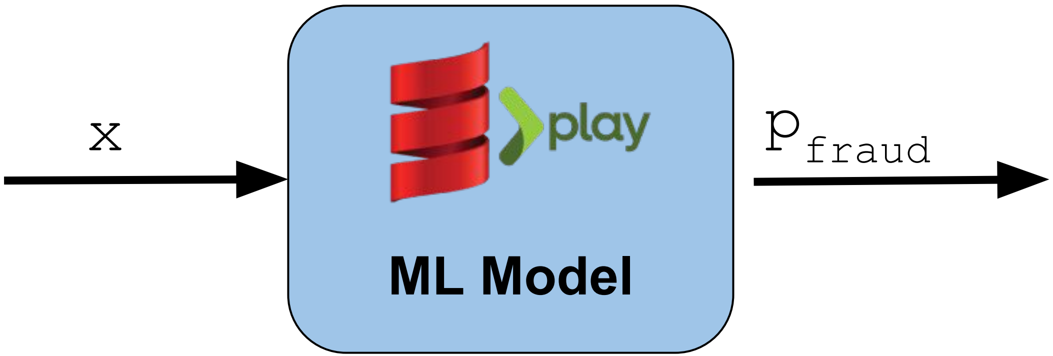

- REST service
- Scala Play service with Spark bindings
- Response time: <1 second

OUR TECH STACK

Scala

amazon web services

GitHub

Spark

docker

EXASOL

R

python

WHY WE SHOULD MONITOR

SCENARIO

Let’s deploy a model for fraud detection in an online shop!

Steps we take:

1. Collect training data.
2. Train a model.
3. Deploy it to production.

COLLECT DATA

Go through the systems and collect data for training

TRAINING DATA

#	Feature-1	Time-to-order [s]	...	Feature-N	**Label**
1	2	300	...	1	**not-fraud**
2	1	5	...	0	**fraud**
3	3	120	...	0	**not-fraud**
4	2	200	...	1	**not-fraud**
5	1	250	...	0	**fraud**
...	...	...	...	...	**...**

FEATURE DISTRIBUTION

Time-to-order [s]
300
5
120
200
250
...

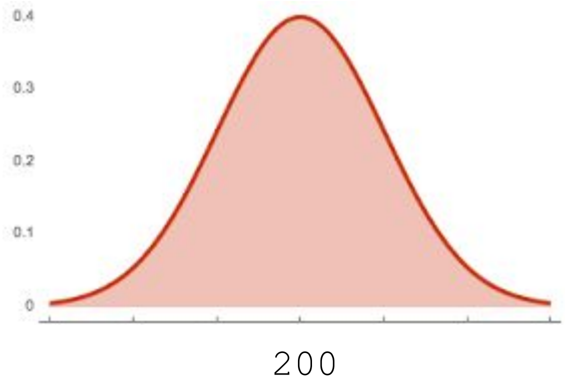

Distribution of feature in training data

GO LIVE

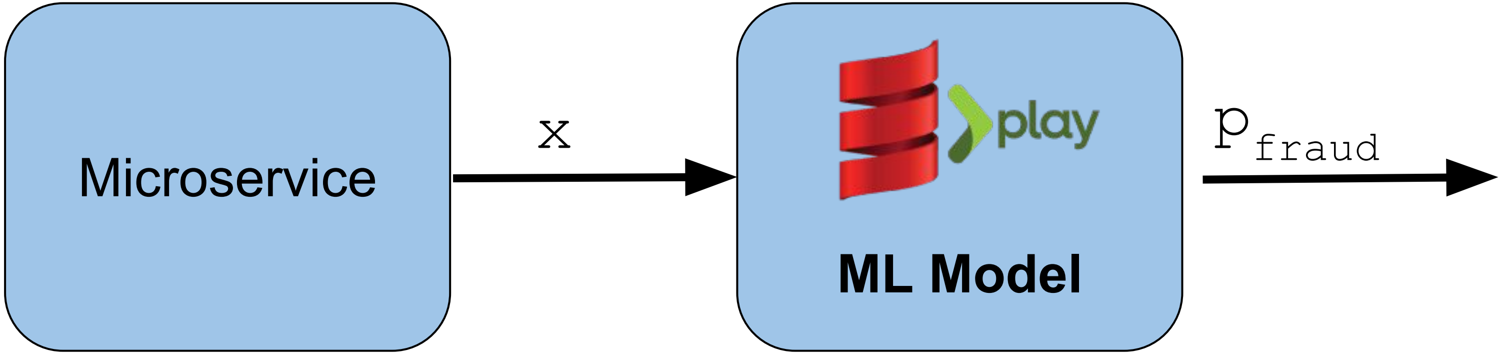

Once we are live, we get features x sent over by a different microservice in real-time.

MONITORING

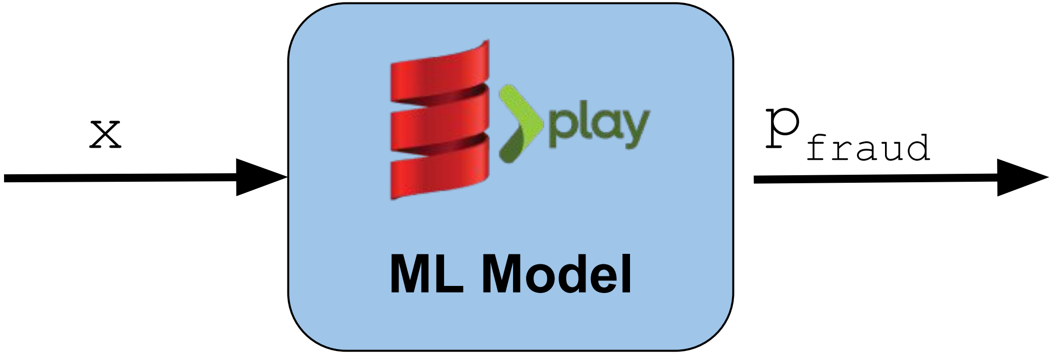

- CPU usage
- Memory usage
- Latency

....

CHANGE OF MOODS

Some weeks later, people are angry:

“We fail to detect fraud, our business is ruined!”

What happened?

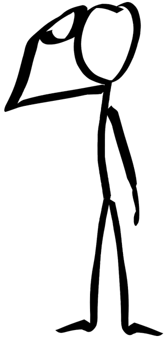

INVESTIGATION

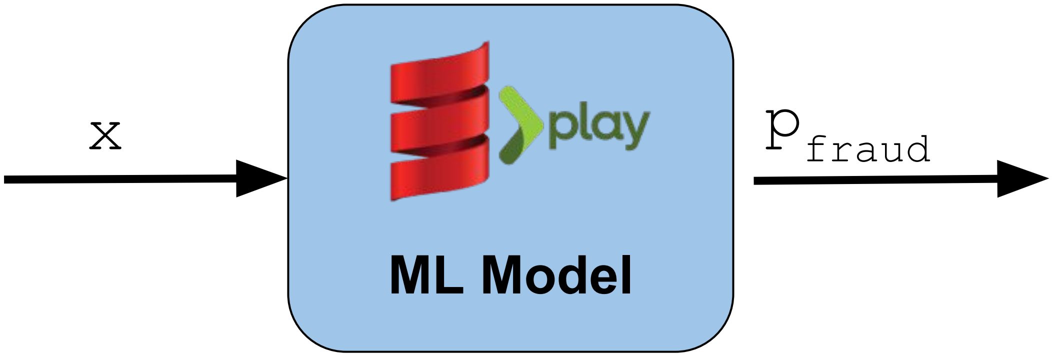

- ✓ CPU usage
- ✓ Memory usage
- ✓ Latency
- ….

INVESTIGATION

ML Model

p_{fraud}

Time-to-order
300000
5000
120000
...

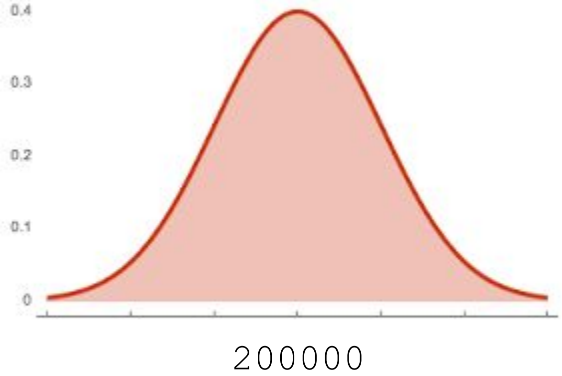

zalando

INVESTIGATION

ML Model

p_{fraud}

Time-to-order
300000
5000
120000
...

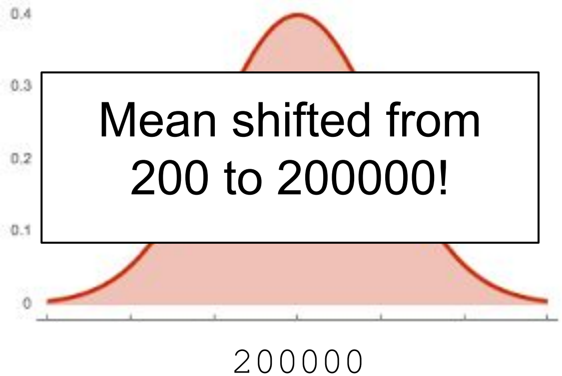

Mean shifted from 200 to 200000!

INVESTIGATION

→ **ML Model** → p_{fraud}

Time-to-order
300000
5000
120000
...

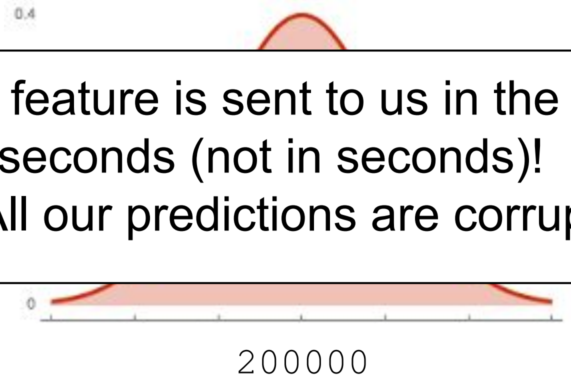

The feature is sent to us in the unit of milliseconds (not in seconds)!
→ All our predictions are corrupt!

PROBLEMS

1. We lost a lot of money.
2. We did not detect it in time.
3. We could have detected it in time and provided a fix.

CONCLUSIONS

We need to make sure that the distributions of input features are (always) the same as in training.

PREDICTION MONITORING

zalando

FAILING FEATURES

Monitor failing input features:

feature name	**fraction**
feature one	**0.903**
feature two	0.004
feature three	0.004
feature four	0.004
...	...

LIVE MONITORING

Compare feature distributions and output probability:

Feature distribution on test data

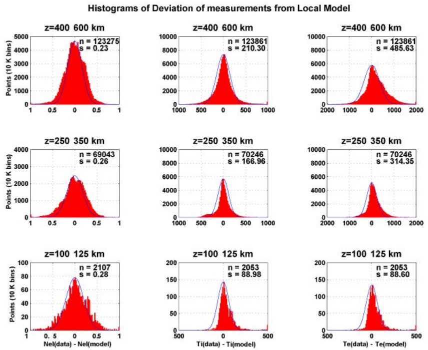

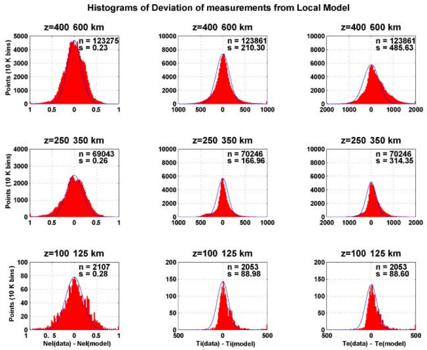

Feature distribution on live data

Quality Monitor

LIVE MONITORING

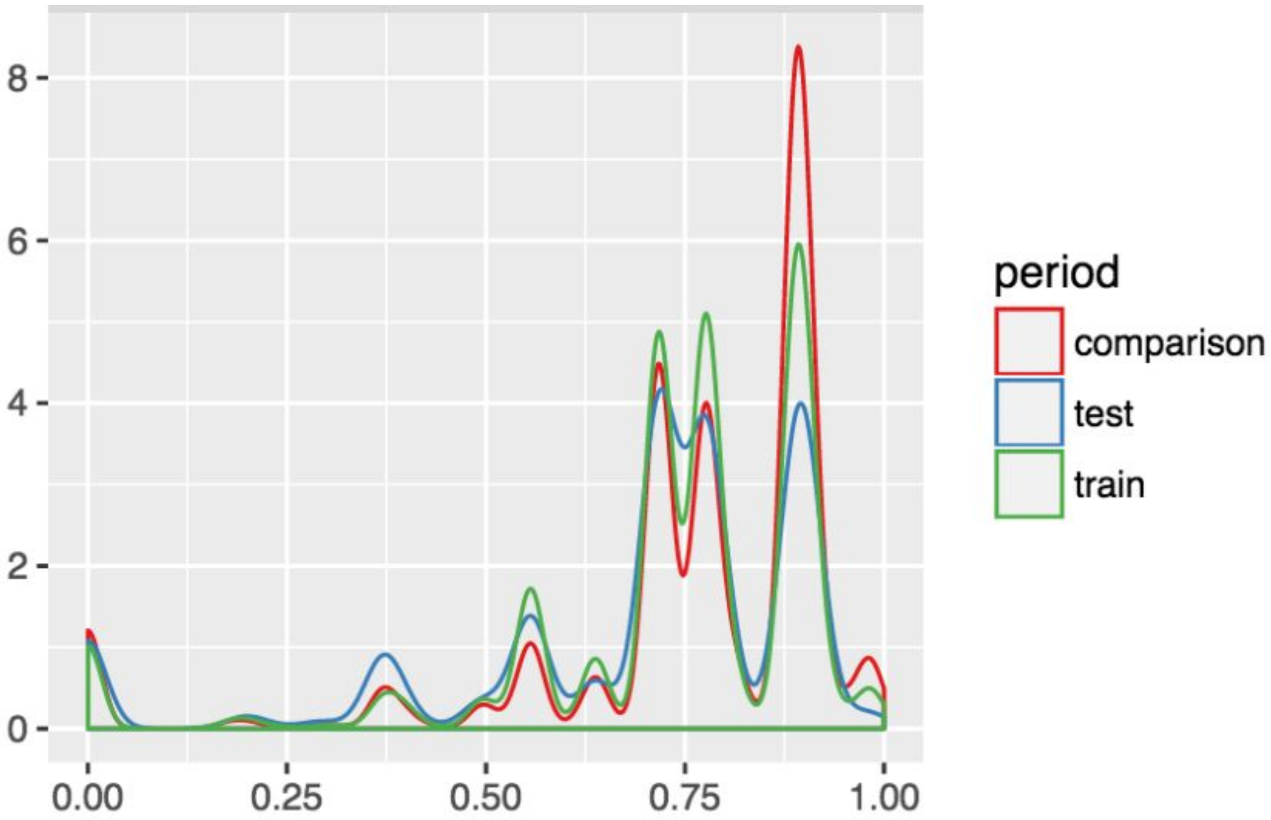

LIVE MONITORING

Compare distributions with KS-distance:

feature name	**this vs previous**	**this vs test**	**previous vs test**
feature one	0.000008	0.928806	0.928798
feature two	0.0009117	0.019504	0.020416
feature three	0.1075305	0.316970	0.313337
feature four	0.943896	0.943655	0.045654
...	...	...	...
prediction	6.606939e-02	0.255182	0.277325

WINDOWS SIZE

How big should the window size for data aggregation be?

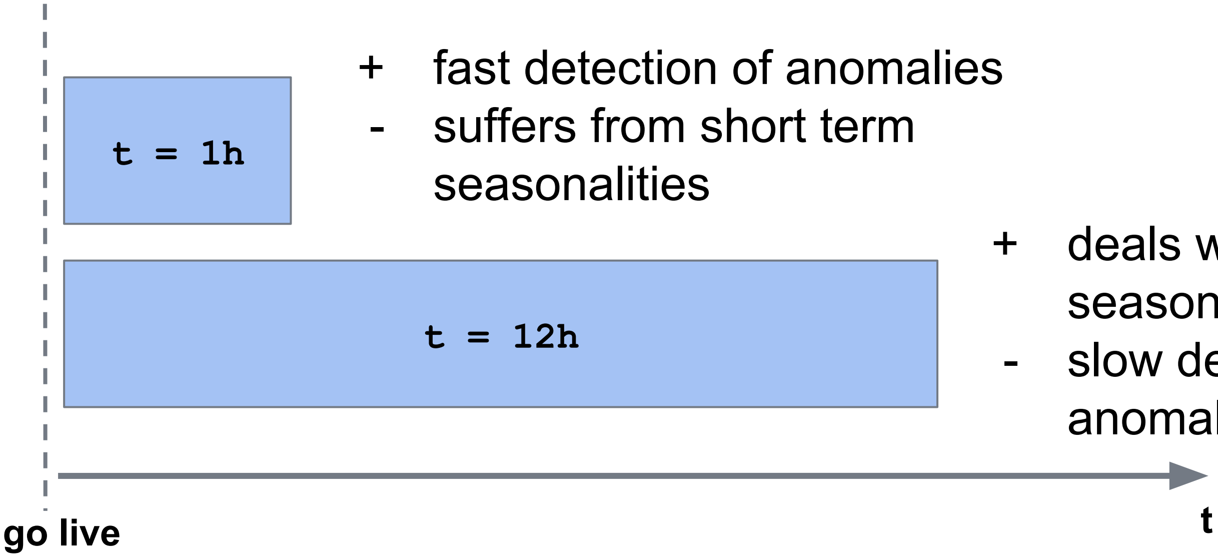

- \+ fast detection of anomalies
- \- suffers from short term seasonalities

- \+ deals with short term seasonalities
- \- slow detection of anomalies

EXECUTION SCHEDULE

How often should we analyze ?

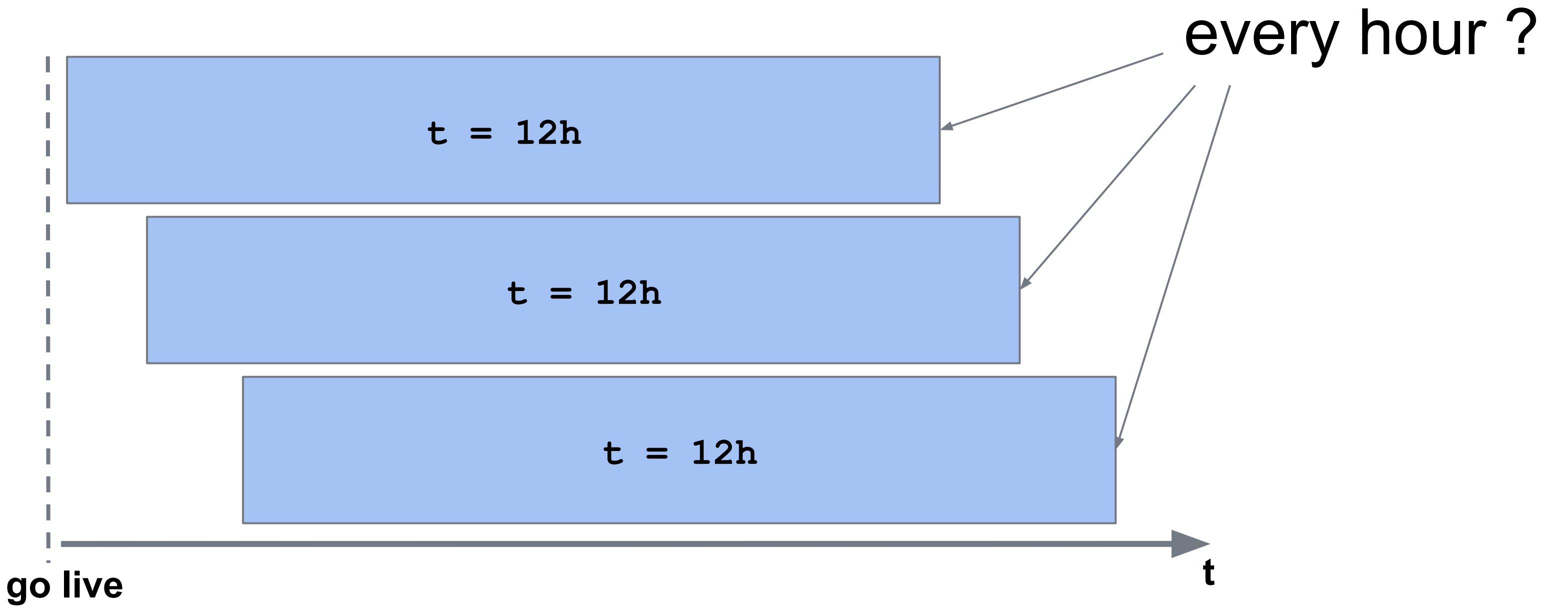

zalando

EXECUTION SCHEDULE

How often should we analyze ?

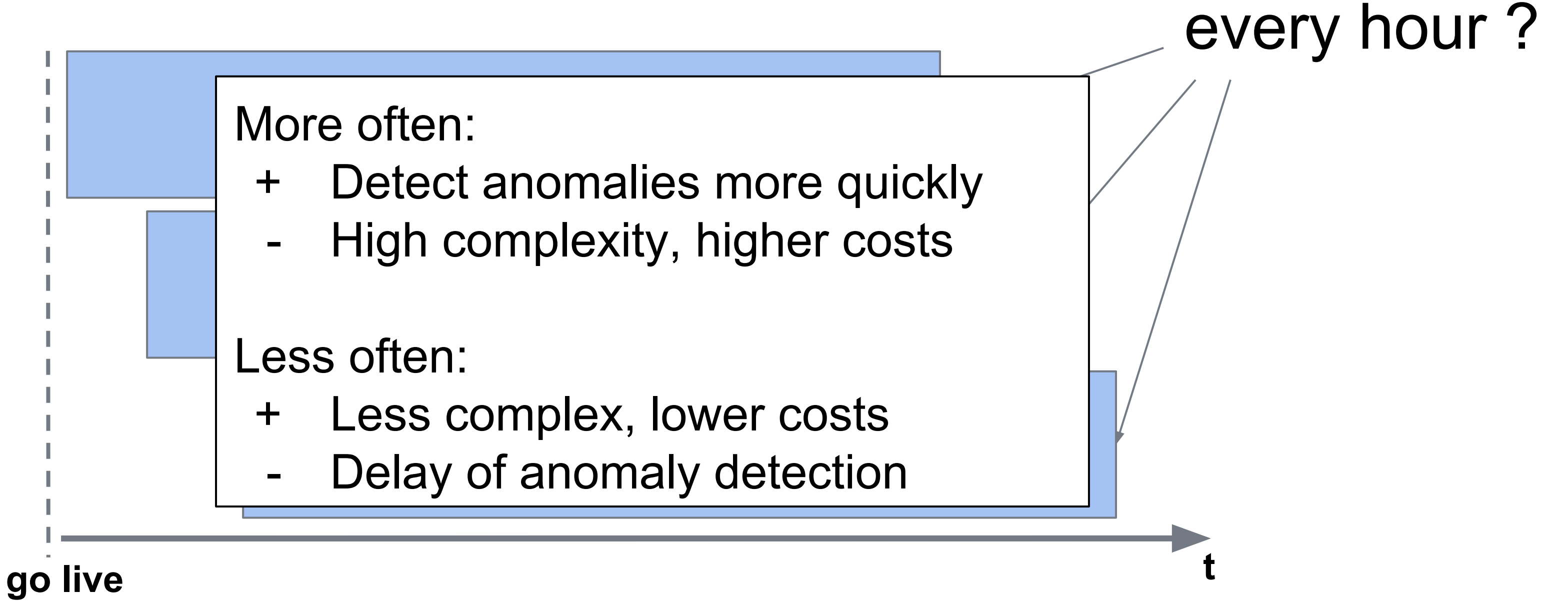

More often:

- \+ Detect anomalies more quickly
- \- High complexity, higher costs

Less often:

- \+ Less complex, lower costs
- \- Delay of anomaly detection

zalando

LIVE MONITORING

possible discoveries

technical problems,
seasonalities,
change of behaviour,
fraud wave,
fraud patterns,
deviation from expectations.

IMPLEMENTATION

DISTANCE BETWEEN TWO DISTRIBUTIONS

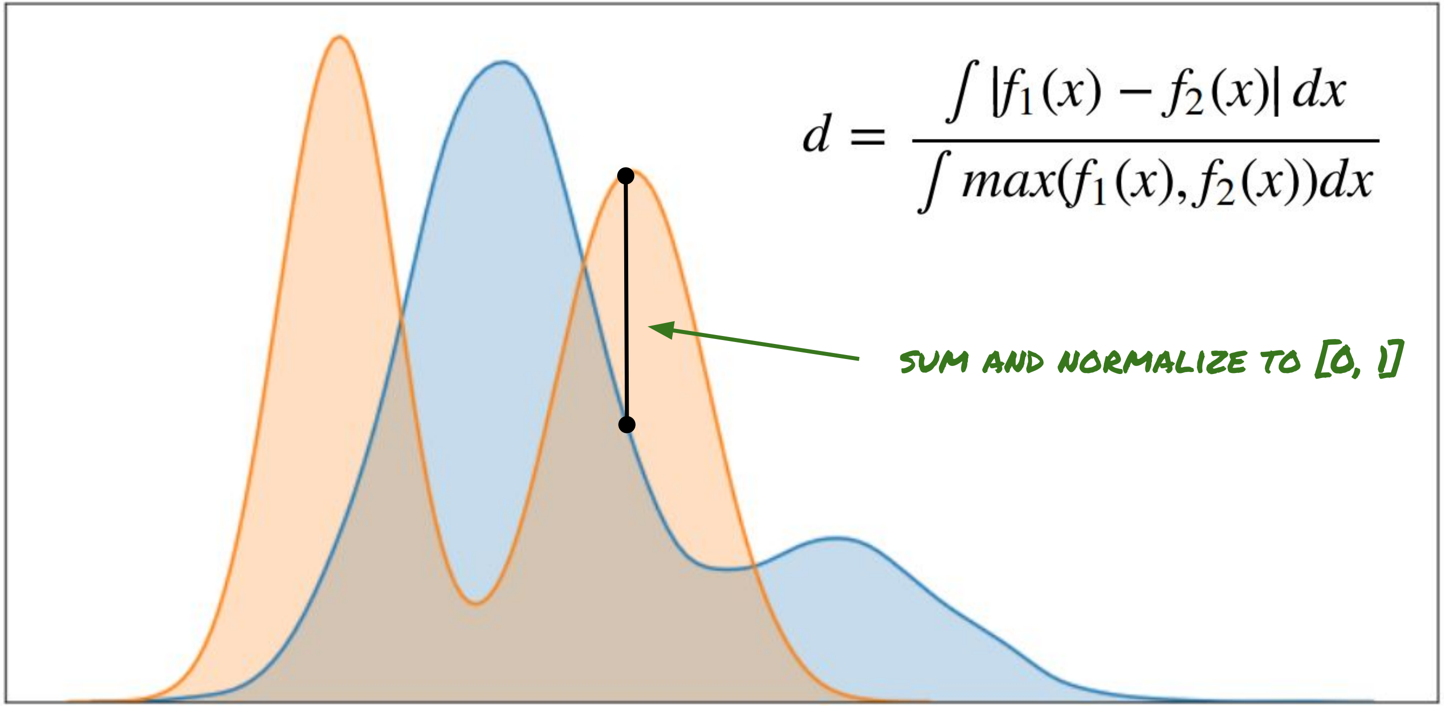

WE USE THE CDF

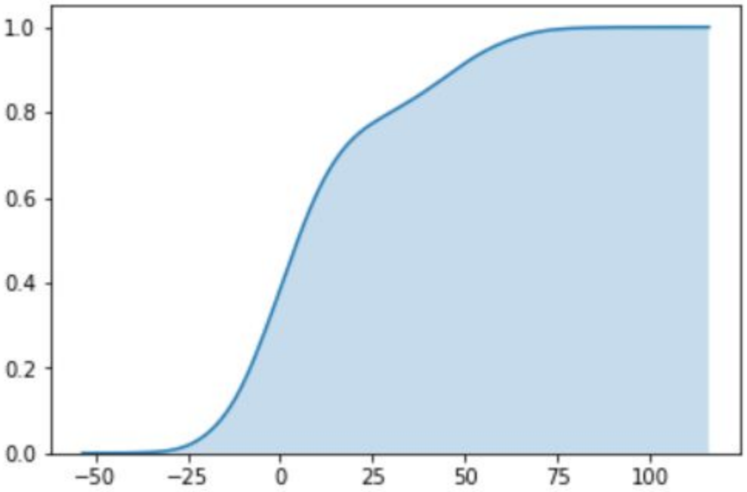

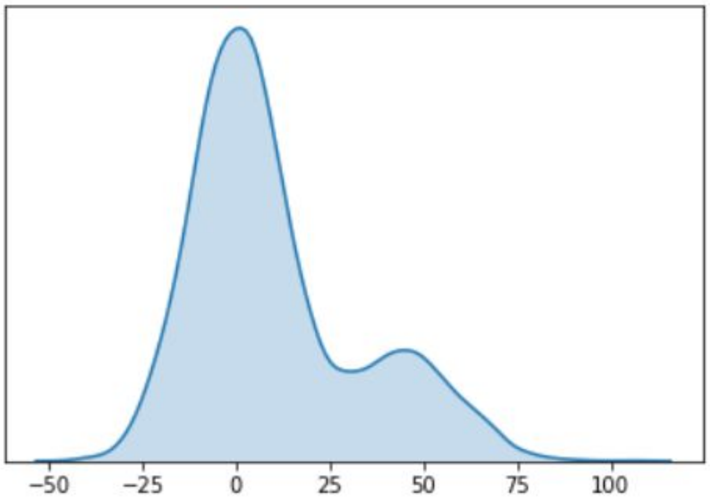

CDF

PERCENTILES

PDF

HISTOGRAM

USING TDIGEST TO OBTAIN CDF

```
import com.tdunning.math.stats.TDigest
import org.apache.spark.rdd.RDD

def create(numbers: Seq[Double]): TDigest = {
    val digest: TDigest = TDigest.createDigest(100)
    numbers.foreach(x => digest.add(x))
    digest
}

def create(numbers: RDD[Double]): TDigest = {
    val empty: TDigest = TDigest.createDigest(100)

    numbers.treeAggregate(empty)(
      seqOp = (acc: TDigest, x: Double) => {
          acc.add(x)
          acc
        },
      combOp = (digest1: TDigest, digest2: TDigest) => {
          digest1.add(digest2)
          digest1
        }
    )
}
```

FROM AN IN-MEMORY COLLECTION

FROM A DISTRIBUTED COLLECTION

zalando

PREDICTION SERVING

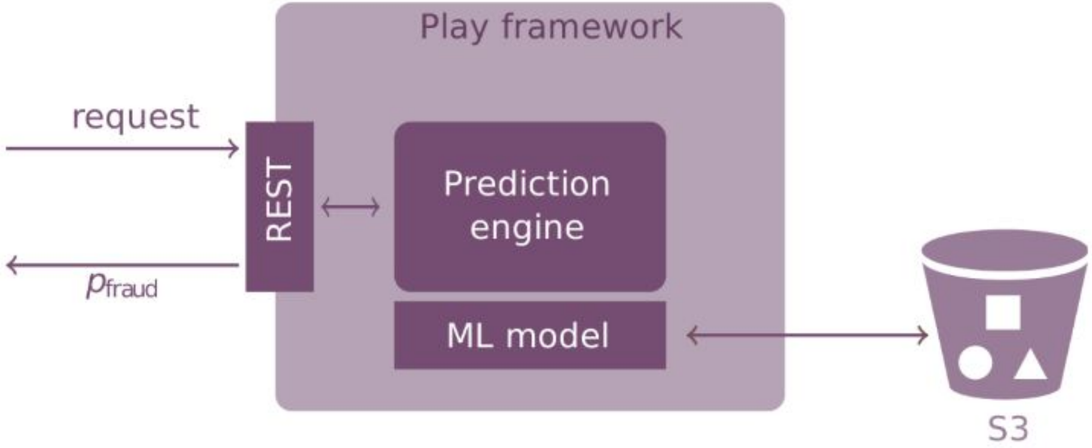

PREDICTION SERVING

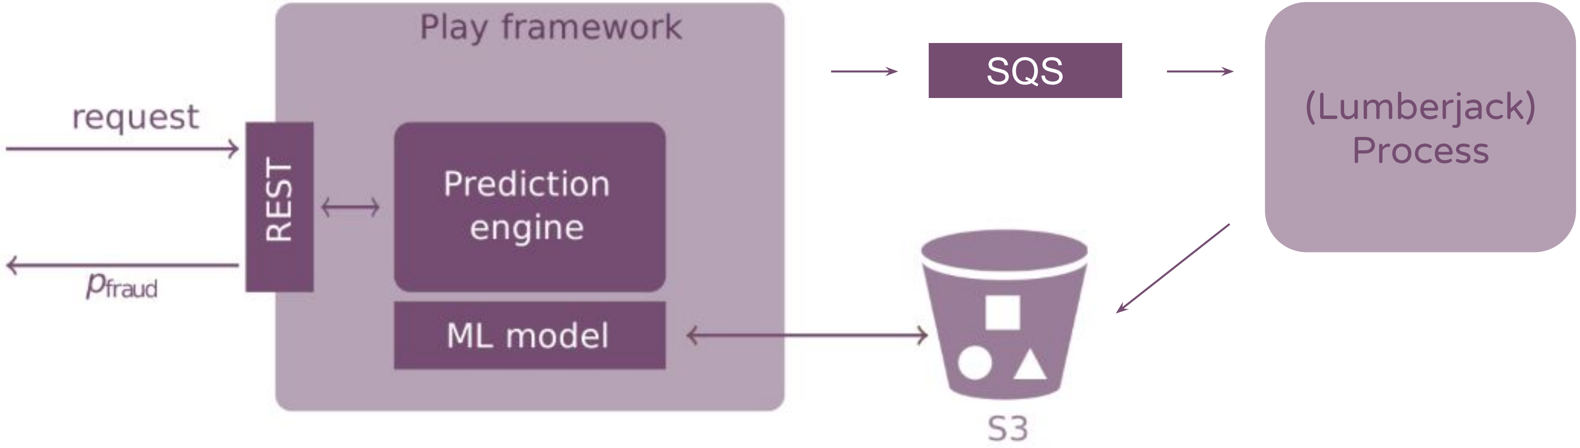

PROCESSING THE SQS MESSAGES

```
func process(sqsClient sqsiface.SQSAPI, dumpSize int,
    interrupt <- chan bool, upload func([]*sqs.Message)) {

        buffer := make([]*sqs.Message, 0, dumpSize)
        timer := time.NewTimer(maxFetchingTime)

        for {
            select {
                case <-interrupt:
                    return
                case <-timer.C:
                    upload(buffer)
                default:
                    for _, message := range receiveMessages(sqsClient) {
                        buffer = append(buffer, message)
                    }
                    if len(buffer) == dumpSize {
                        upload(buffer)
                    }
            }
        }
    }
```

PUTTING IT TOGETHER IN AWS DATA PIPELINE

logs

- process logs
- group by models
- get failed features
- create tdigests
- create histograms
- feature distances

Reports

Alerts

FINAL NOTES

- if you have a ML system deployed in production, then you have to monitor it somehow
- monitoring is especially important if performance feedback is delayed
- start simple and non-intrusive, keep the reports flexible
- automate as much as possible
- to measure how far you are with monitoring, go through the questions in this paper from Google: *"What’s your ML Test Score? A rubric for ML production systems"*

THANK YOU!

Patrick Baier & Lorand Dali

https://tech.zalando.com/blog/scalable-fraud-detection-fashion-platform