LESSONS LEARNED WITH MLOPS

ATTILA TOZSER

03.05.2022



AGENDA

INTRODUCTION

ML COMPLEXITY

USE-CASE 1: UTILITIES

USE-CASE 2: MANUFACTURING

INTRODUCTION

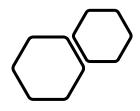
Attila Tozser

- 10+ years of Data Engineering and 5+ years Cloud experience
- 20+ ML projects different sizes.
- Working on the German market with customers usually PB range data size and several Data Science teams.
- Consulting for on premise to cloud migration and creating products out of data science ideas

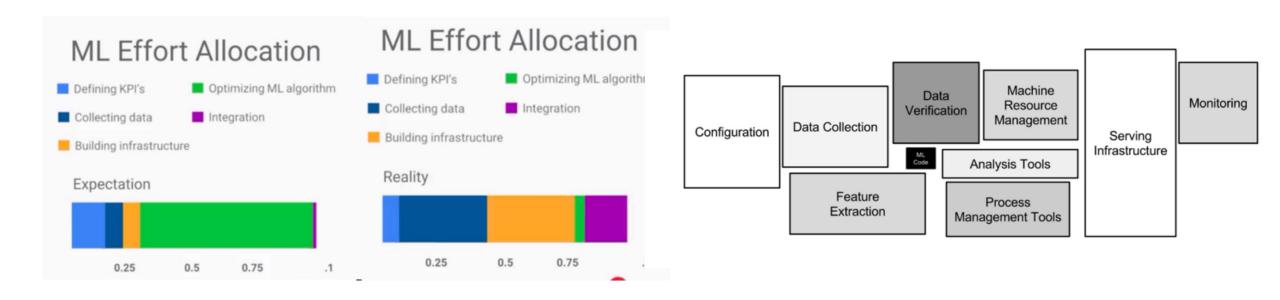




ML COMPLEXITY



WHERE THE EFFORT GOES?



Hidden Technical Debt in Machine Learning Systems (neurips.cc)

https://www.coursera.org/learn/end-to-end-ml-tensorflow-gcp/home/welcome

WHERE THE EFFORT GOES?

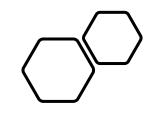
Role	Workload
Sotfware Engineer, Machine Learning Engineer, Data Engineer	0.7
Data Scientist, Machine Learning Researcher	0.2
Al product Manager	0.1

pz 5pz #ccc}.gbrtl .gbm{-moz-be #ecc;display:block;position:absol opacity:1;*top:-2px;*left:-5p* /stop:-4px\0/;left:-6px\0/ me-box;display:inline-block implay:block;list-style:none lock; line-height: 27px er;display:block; 1000).0



UTILITIES





7

CC). dDrt1 lay:block;position:a 50k employees DAX company Would like to have AI and ML capability. Create a department from scratch. (2017-2019) 8 to 50 heads within two years

BUSINESS ENVIRONMENT

Complexity meets reality (how it turned out in reality)?

Sotfware Engineer Machine Learning Engineer

Data Engineer

Data Scientist

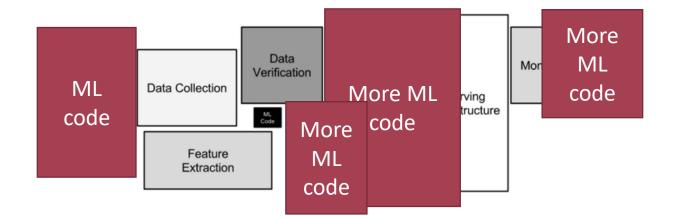
Machine Learning Researcher Al product Manager

The complexity is still not machine learning.

What use-case to prioritize?

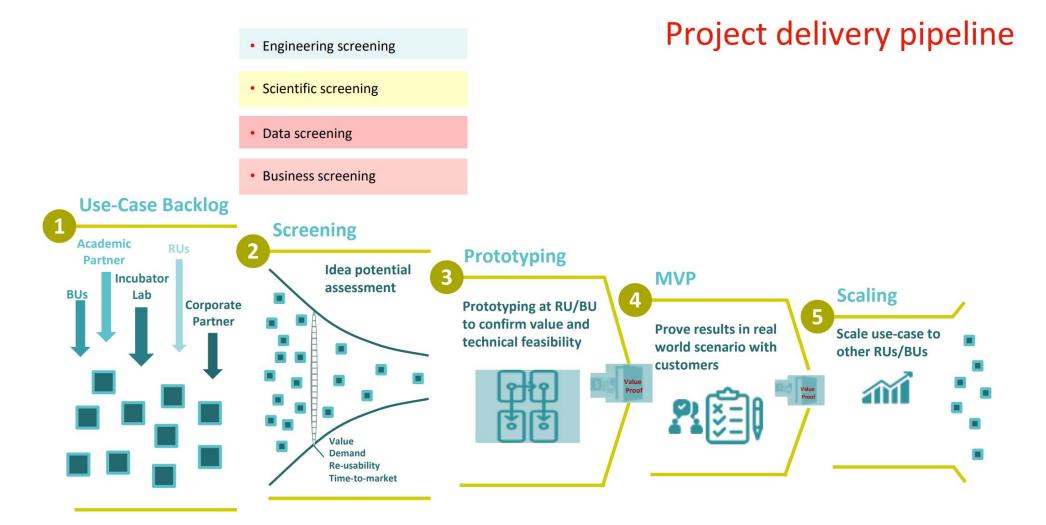
Where to get knowledgeable people? How we setup an application?

Growth 7 + 1 -> 50 + 10



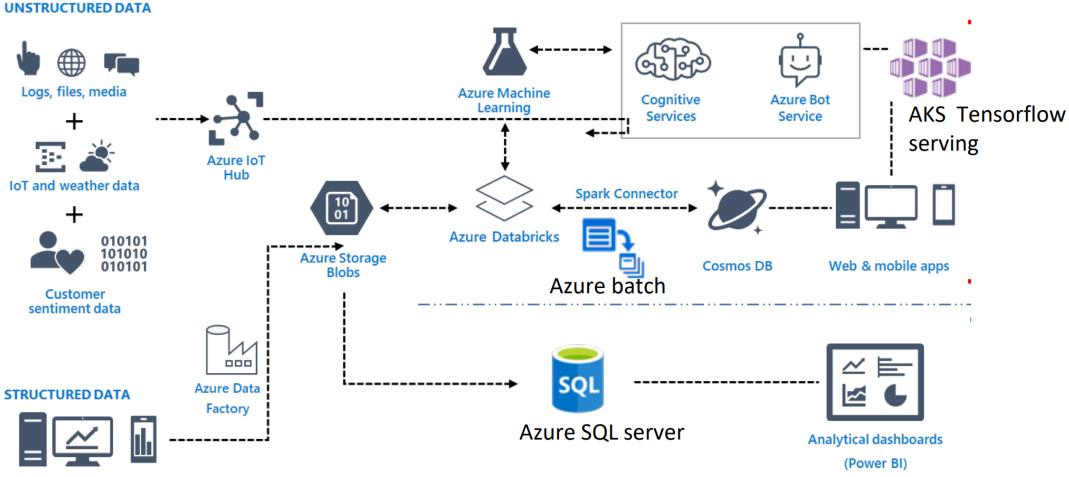
How we qualify cases?

Large company with a bit of internal advertisement generated ~200 prospects within 3 month.



#ccc}.gprt1 .g isplay:block;position:a Cases are very different in terms of technology, data size, used models, business requirements and latency needs lock;list-sty ne-height

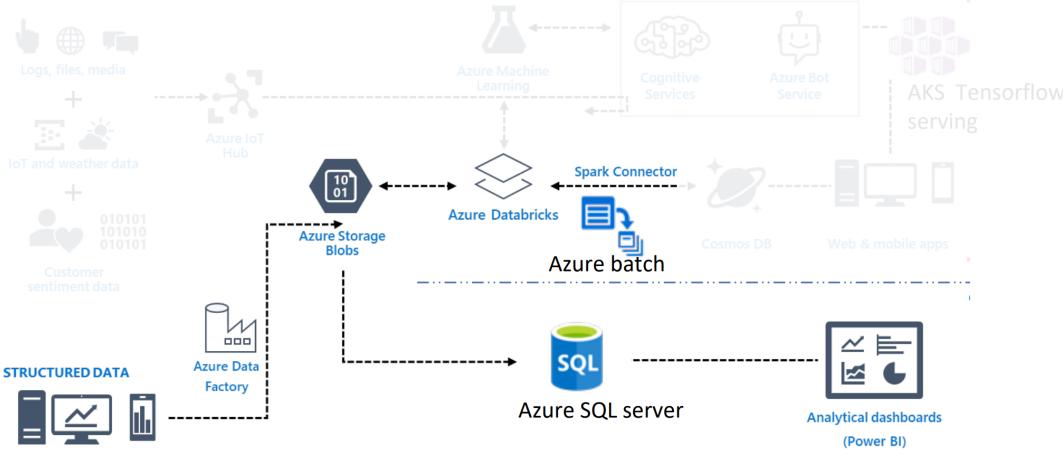
1 SIZE FITS ALL ARCHITECTURE (AZURE IN 2019)



Business apps

1 SIZE FITS ALL ARCHITECTURE (AZURE IN 2019) NON-REALTIME EXAMPLE

UNSTRUCTURED DATA



Business apps

Bayernwerk entwickelt Energie-Monitor für Kommunen: Zeitung für kommunale Wirtschaft (zfk.de)

Der EnergieMonitor als innovatives Tool für die Energiewende [Werbung] (trendsderzukunft.de)

Realtime renewable energy production data on dashboards in the town hall.

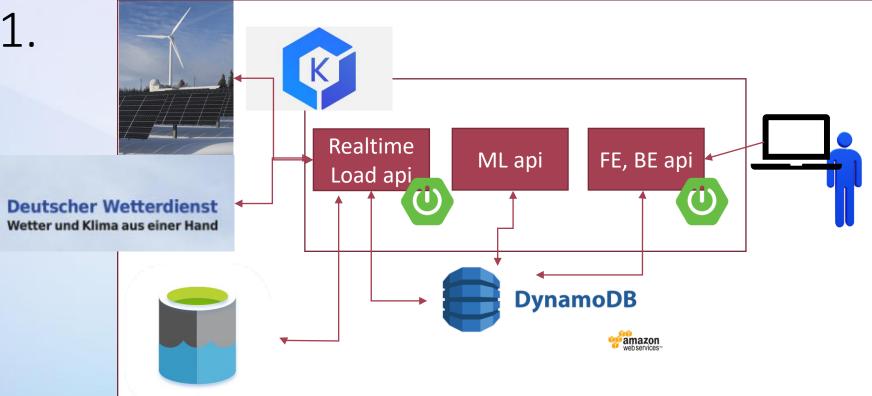
USE CASE 1



TECH ARCHITECTURE NO 1.

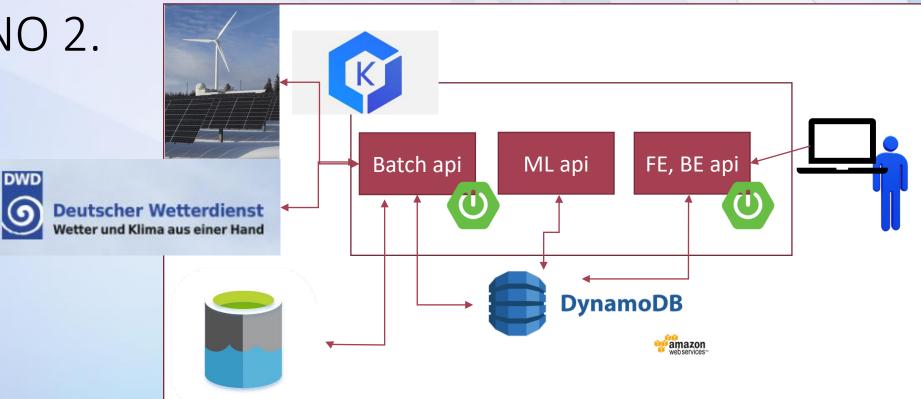
DWD

- Realtime production data from the Power plants
- Realtime weather data from the german weather services
- Java spring boot apis to load the content to Dynamo \rightarrow trigger ML
- Kubernetes ML model serving.
- Real-time api to read Dynamo and update the dashboards in the town hall.
- Sounds good?



TECH ARCHITECTURE NO 2.

- 'Realtime' production data from the Power plants (daily csv)
- Realtime weather data from the german weather services
- Java spring boot apis to load the content to Dynamo → trigger ML
- Kubernetes ML model serving.
- Real-time api to read dynamo and update the dashboards in the town hall.
- Sounds good?

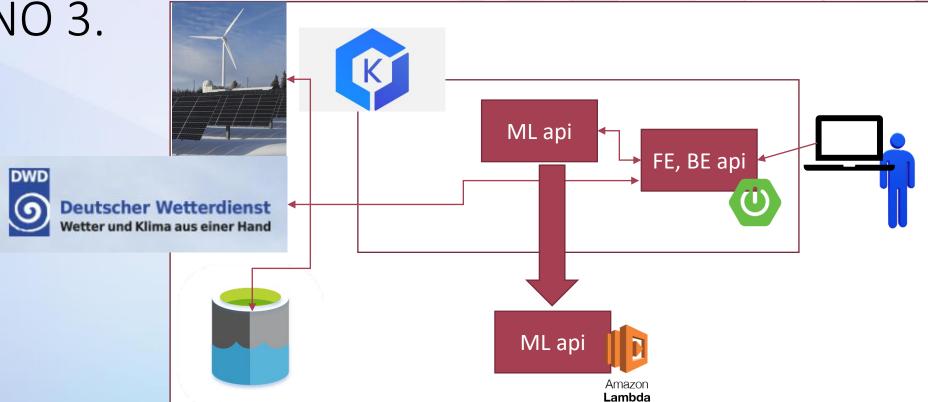


• Simple physical model is predicting it with 85% accuracy with weather data

• Consumption is stable (best model is last 10 days average, with differentiating the weekend)

• Production architecture is very stable not changing over time not so many in one municipality (<u>only maintenance needs consideration</u>)

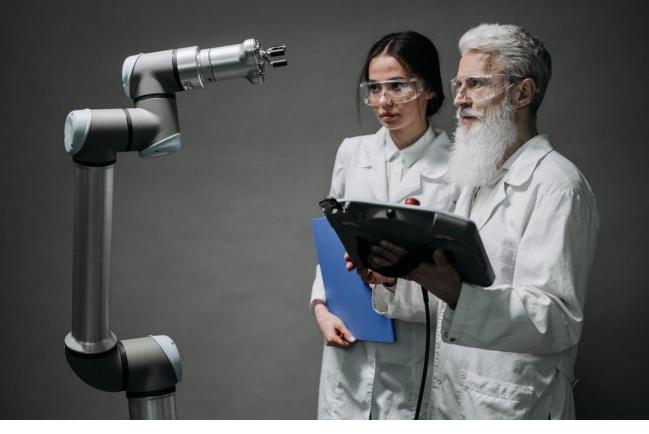
• UI, calling weather services api, calling AWS lambda and showing the results.



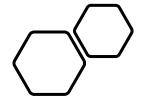
TECH ARCHITECTURE NO 3.

px #ccc}.gprt1 .gpm(-moz-D display:block;position:absol opacity:1;*top:-2px;*left:-5; Revisiting the architecture saved 95% of the costs for the ML part. lepley:block;list-style:none; nemblock:line-height:27px

Spx #ccc}.gbrtl .gbm{-moz-be display:block;position:absolu icity:1;*top:-2px;*left:-5px top:-4px\0/;left:-6px\0/;ri -box;display:inline-block;fc splsy:block;list-style:none; -block; line-height: 27px; padd ointer;display:block;text ive;z-index:1000).gbts/



MANUFACTURING



USE-CASE 2

y:block:pos Consumer goods company (10k + employees), central IT Data organization is responsible for Data Science and Engineering with expert teams. Mature tools \rightarrow full ML stack on Azure with multiple options CUSTOMER

USE-CASE: SHEET BREAK ANALYSIS FOR PAPER MILLS

(PDF) Detection, Diagnosis and Root Cause Analysis of Sheet-Break in a Pulp and Paper Mill with Economic

Impact Analysis (researchgate.net)





Strength Machine Direction (MD)



Wet Strength Machine Direction (MD)

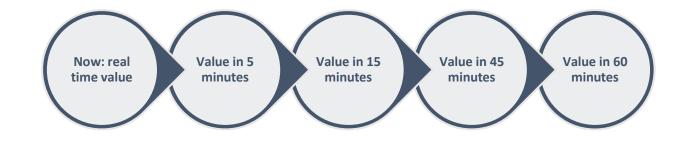


USE-CASE: SHEET BREAK ANALYSIS FOR PAPER MILLS

Real time sampling of machine parameters fed into machine learning algorithms

This is a sheet-break, causes:

- 2 min lost time in Mill
- 40 min lost time in Cleaning
- Losses in energy, raw material...

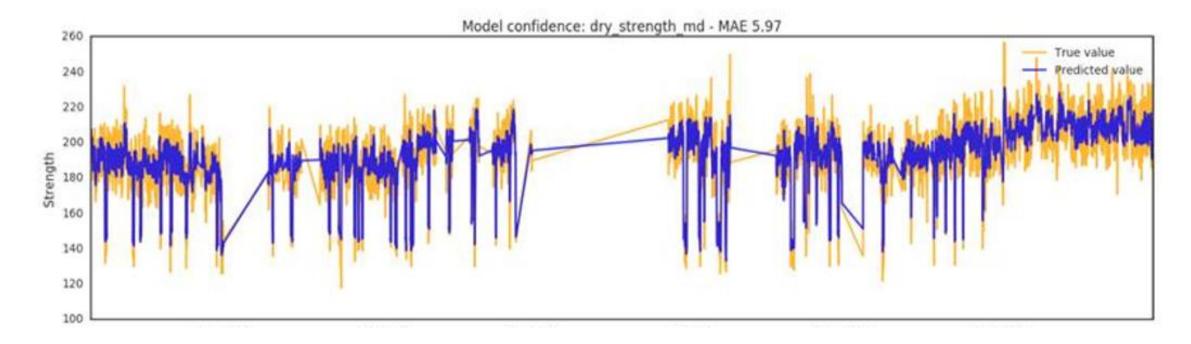


After	Before
1 prediction per minute.	1 measurement every 80 minutes
Prediction without lag time: the value is predicted just at the moment.	Measurement obtained 20 min after the sample is produced
Prediction is kind of averaged value, better for process control	Lot of noise in the measurement

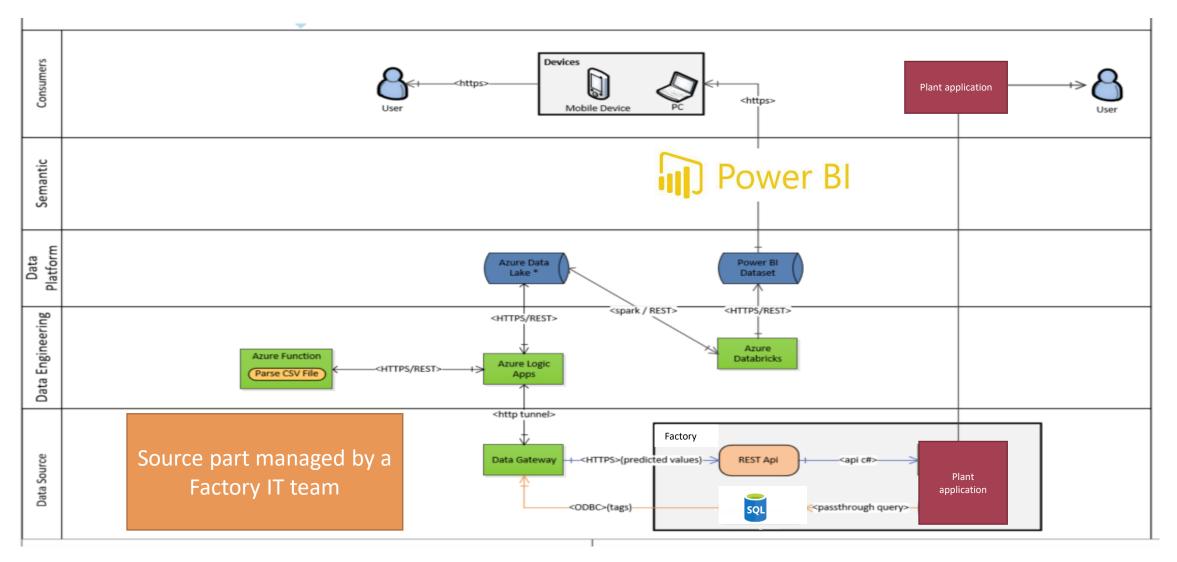
THE MODEL:

The model delivered by an external expert company.

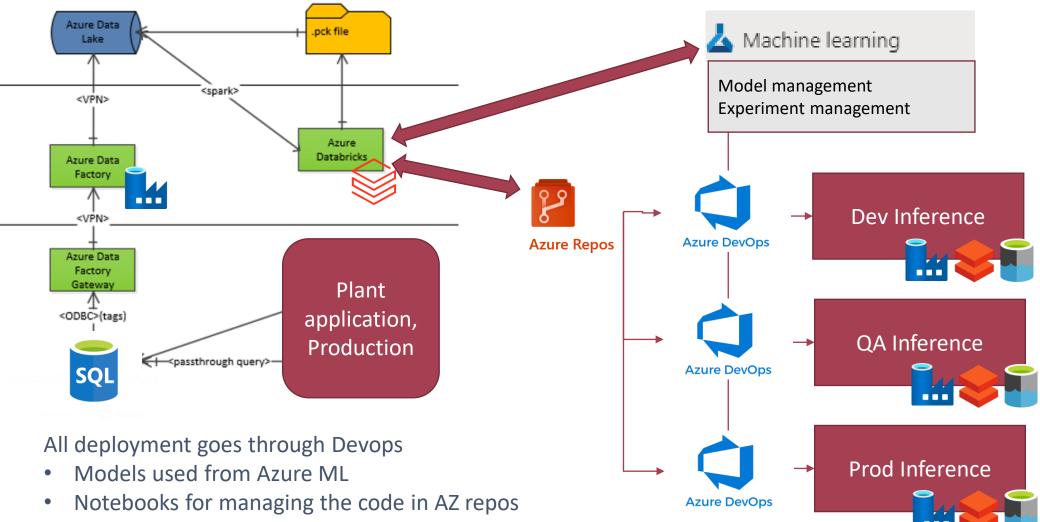
They are Good on papermills, but bad in Software development. It is a notebook 😕



TWO PIPELINES (INFERENCE):



TWO PIPELINES (TRAINING):



• All infra code in the repos, and automation in Devops pipelines

splay:block;positio Difficult to quantify, but considered as success and initiated several additional projects. Due to the tools maturity the deployment and re-training of the models is fully automated, additional factories can be connected with minimal efforts

RESULT



SUMMARY

TOOLS EVOLVED A LOT OVER THE PAST YEARS

Regularly revisit your architecture

DIFFERENT ORGANIZATIONAL MODELS CAN WORK WELL

It all depends on your needs: expert team, cross functional team

DO NOT FEAR THE SIMPLICITY

No BIG data or no huge models maybe not a problem

THANK YOU



ATTILA TOZSER



attila.tozser@industrialize.ai

WWW.INDUSTRIALIZE.AI (coming soon)