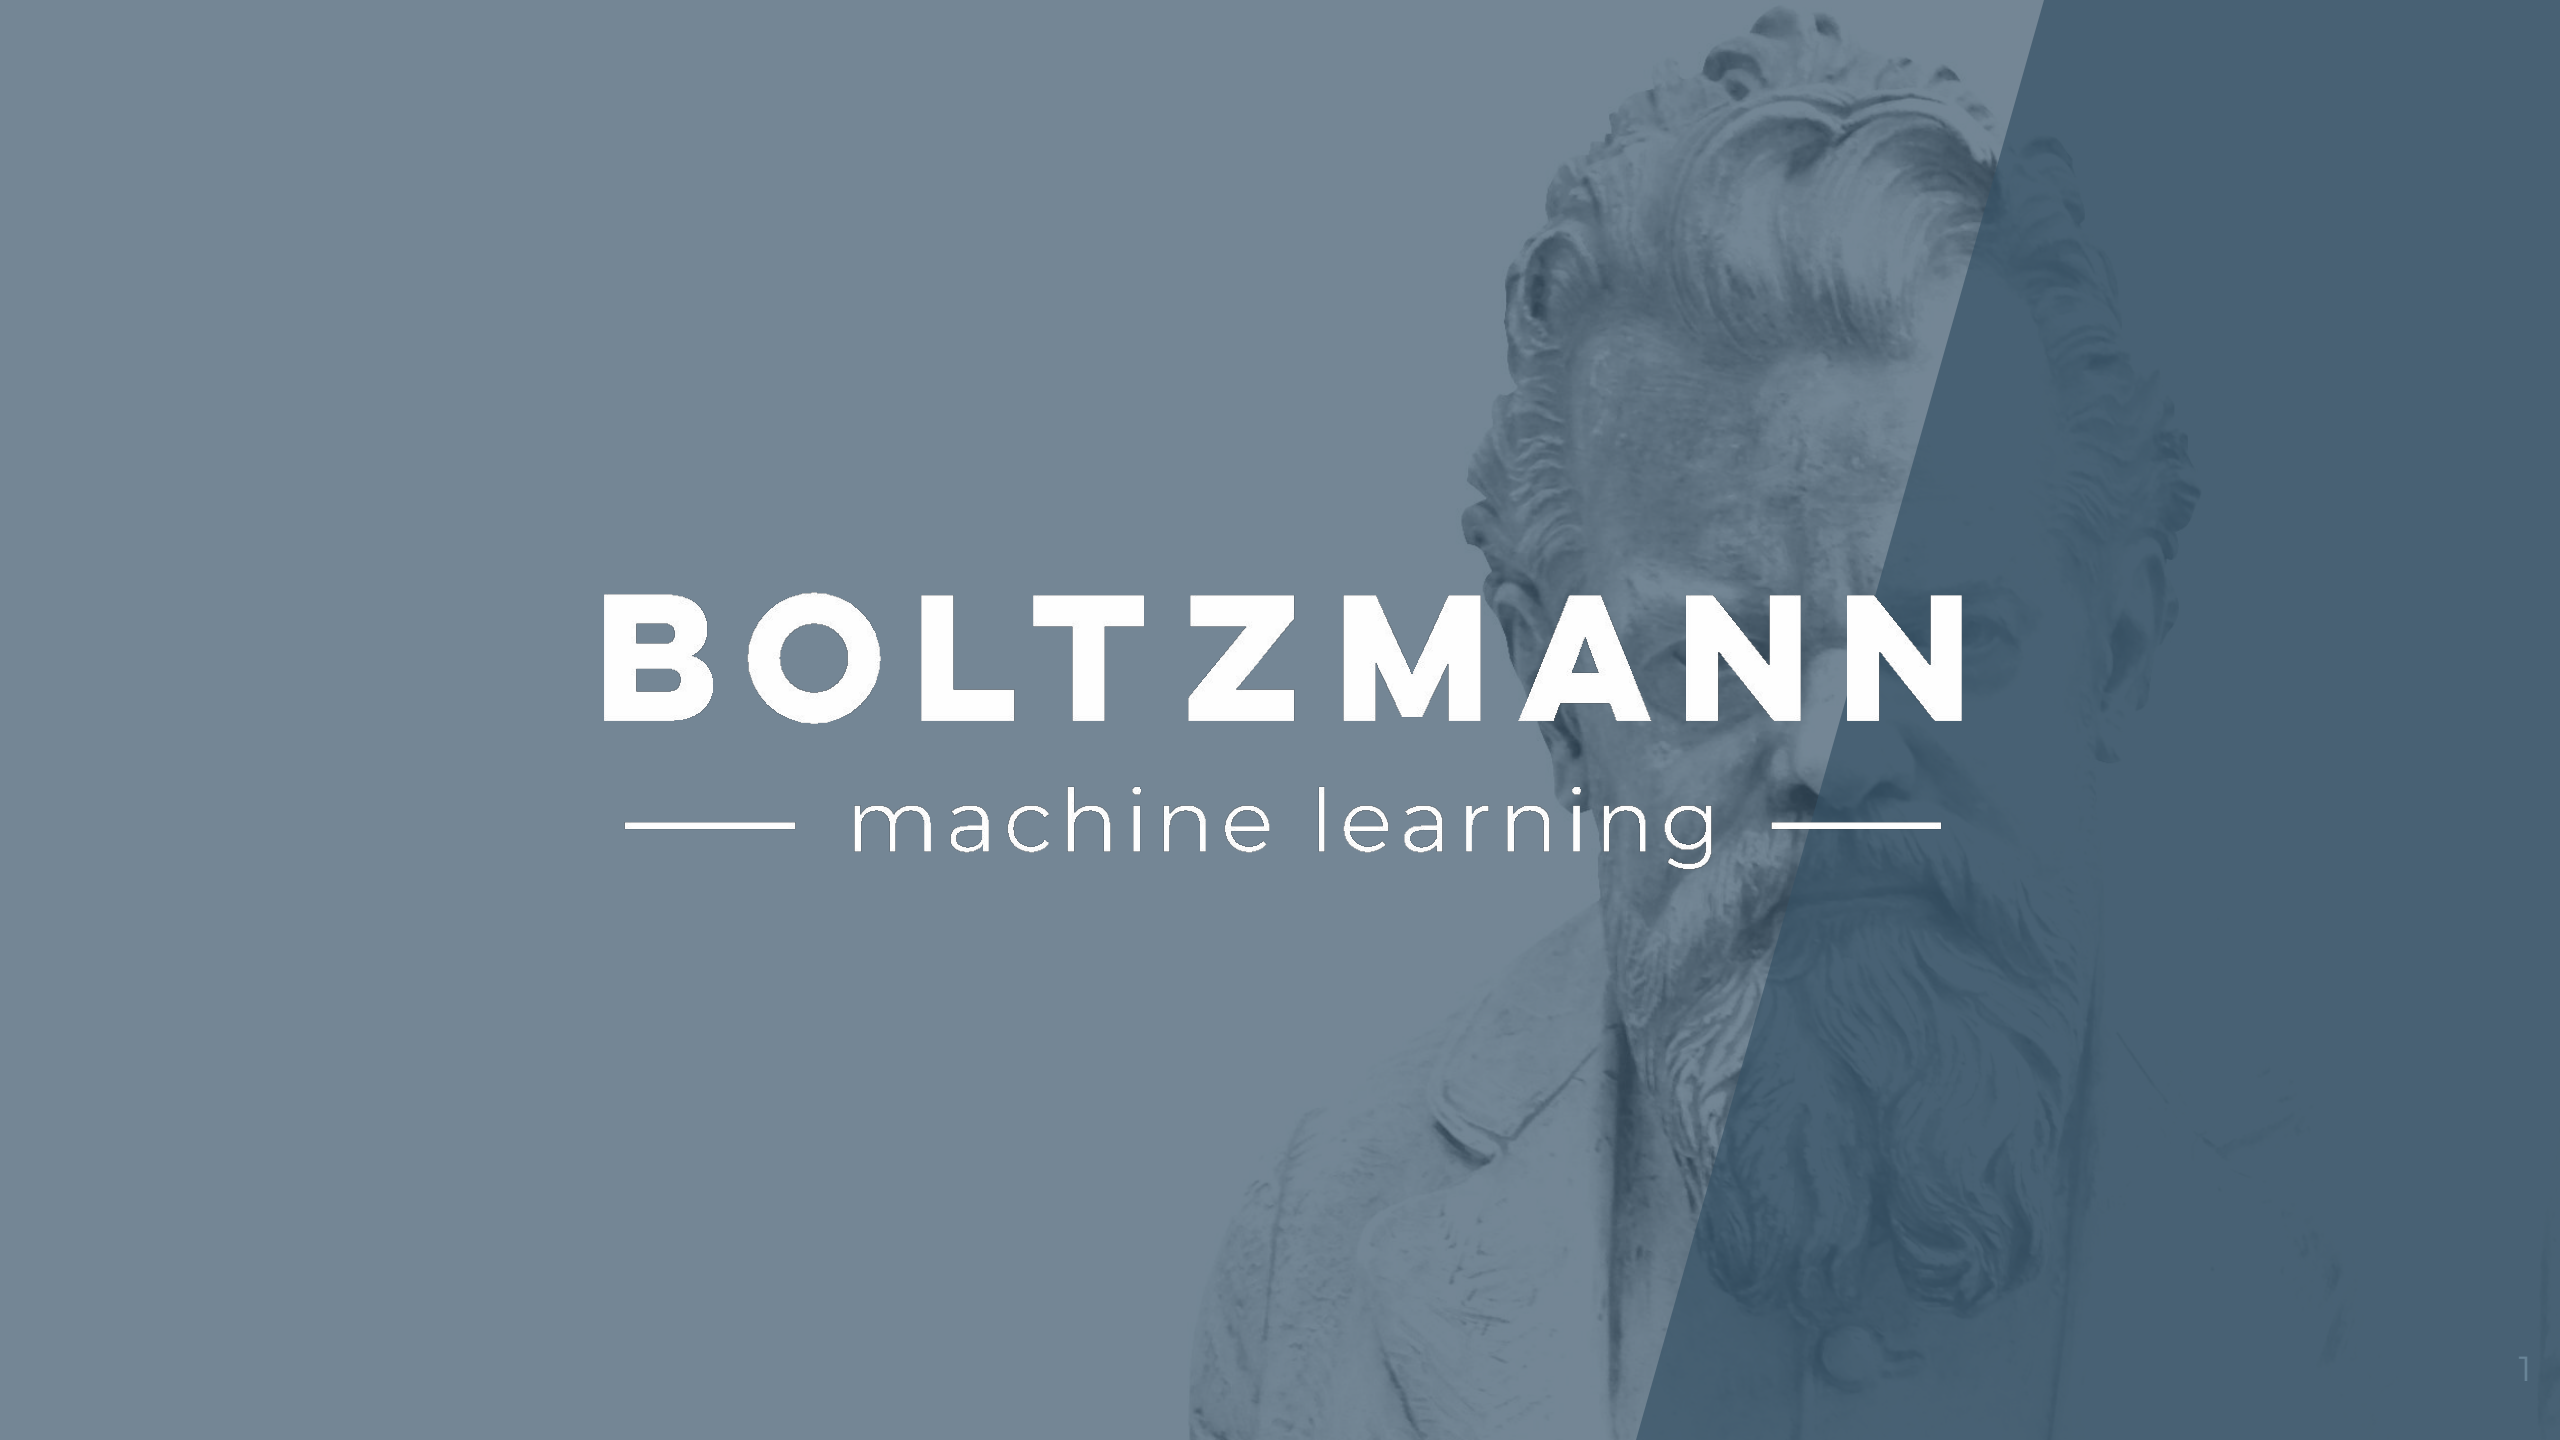
BOLTZMANN
machine learning

Boltzmann

- Founded in 2017 by AI specialists
- Team of 15 incl. 2 professors, 6 PhDs, professional developers & accountants
- Activities:
  - Large clients in consulting (banks, energy, industry, ...)
  - **Ludwig Assistant**, the AI assistant for accountants. First pilot customers in 2018, now more than 15.000 companies

BOLTZMANN

machine learning

Projects

LUDWIG

BASF
We create chemistry

luminus

AUCXIS

Co-libry

contractify

TOTAL

EY

Deloitte.

Bishop Fleming

titeca accountancy

SOME OF OUR WORK

Fraud detection, AML, Invoice mgt prediction

@ Finance

Predictive Maintenance, Utilities prediction

@ Energy

Contract clauses, Patents discovery, Product categorisation

@Data platforms, Retail

Cold start recommender

@ Immo

Wage benchmark

@ HR

LUDWIG

@ Accountancy

We build
MACHINE LEARNING
applications for the
CORE PROCESSES
of your company

1 Co-discover and **IDENTIFY** the best ML application

2 **TOGETHER** develop a tracer-bullet MVP

3 **INDUSTRIALIZE** the implementation

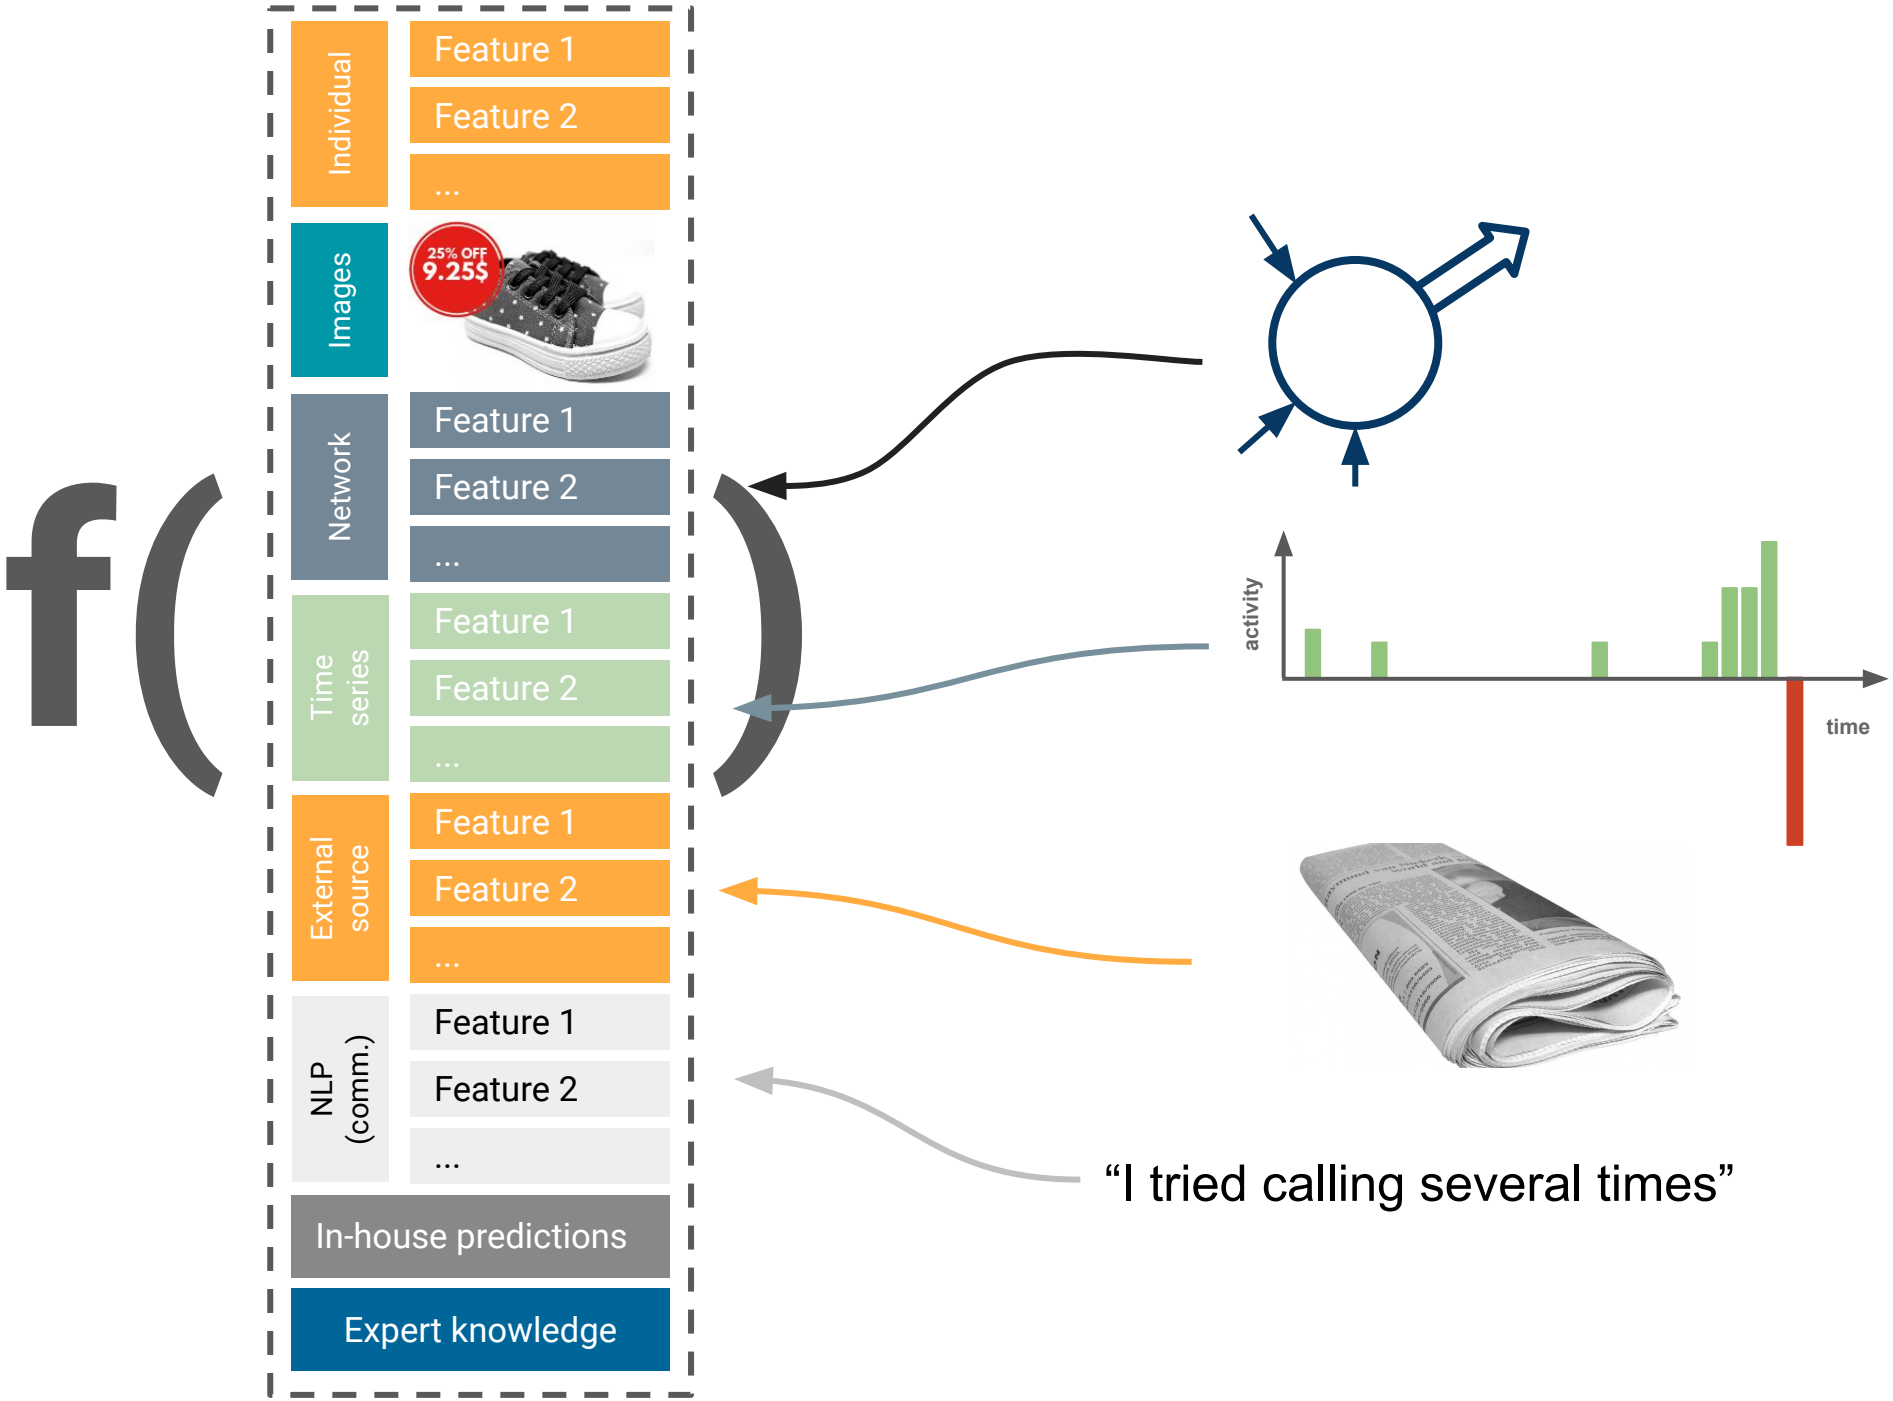
Individual
Feature 1
Feature 2
...
Images
25% OFF
9.25$
Network
Feature 1
Feature 2
...
Time series
Feature 1
Feature 2
...
External source
Feature 1
Feature 2
...
NLP (comm.)
Feature 1
Feature 2
...
In-house predictions
Expert knowledge
f(
)
activity
time
"I tried calling several times"

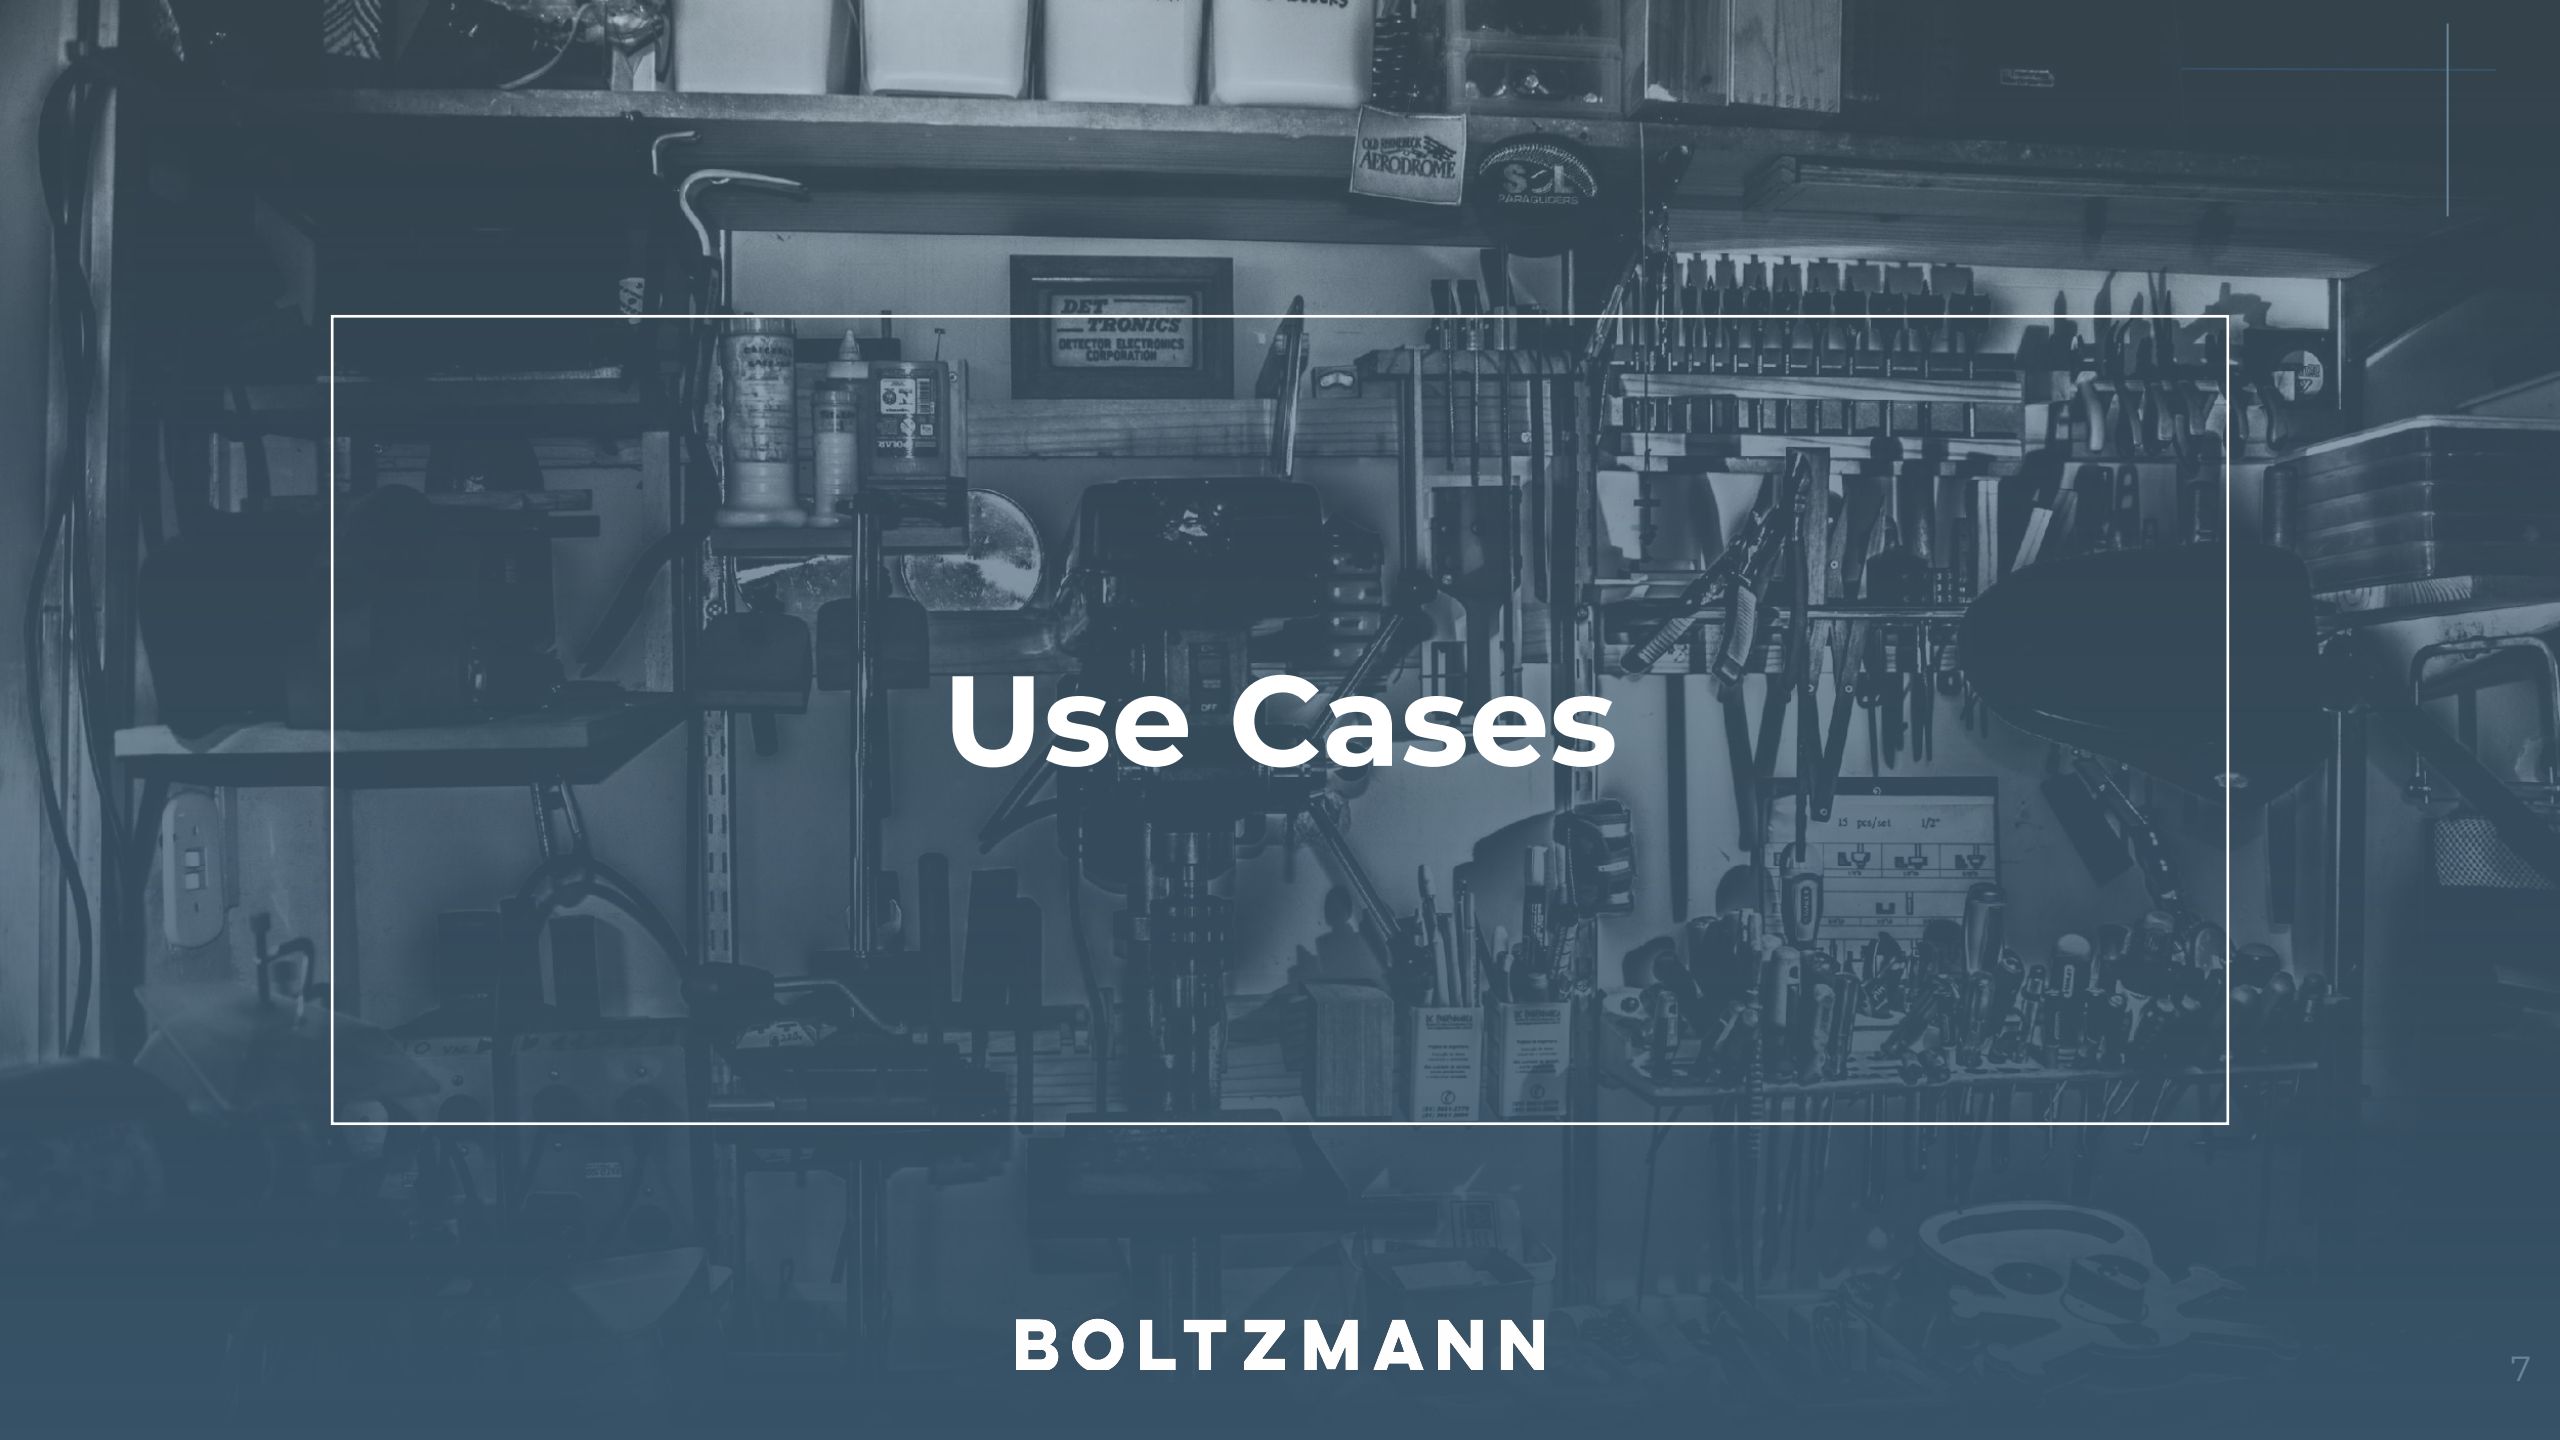

Use Cases

BOLTZMANN

BASF | Predicting utility consumption

Description

- Utility dpt is responsible for **planning utilities** (electricity, gas, steam, etc)
- There are ~**250 planning points** in total spread across multiple sub plants
- Available data sources:
  - **(Chemical) production** planning and actuals
  - **Historic** consumption / production
- Limited to **no time series characteristics**:
  - No autocorrelation, seasonality, …
  - Biggest predictor: production planning

Challenges and solutions

- Different characteristics across sub plants and planning points
  - Custom (lightweight) **AutoML**
- Deployment: complicated architecture with many on-premise servers and security layers
  - **Azure implementation:**
    - Azure Data Factory
    - Azure ML Pipeline
    - Azure Function
- Confidence in prediction
  - Many tailored **'uncertainty' flags**

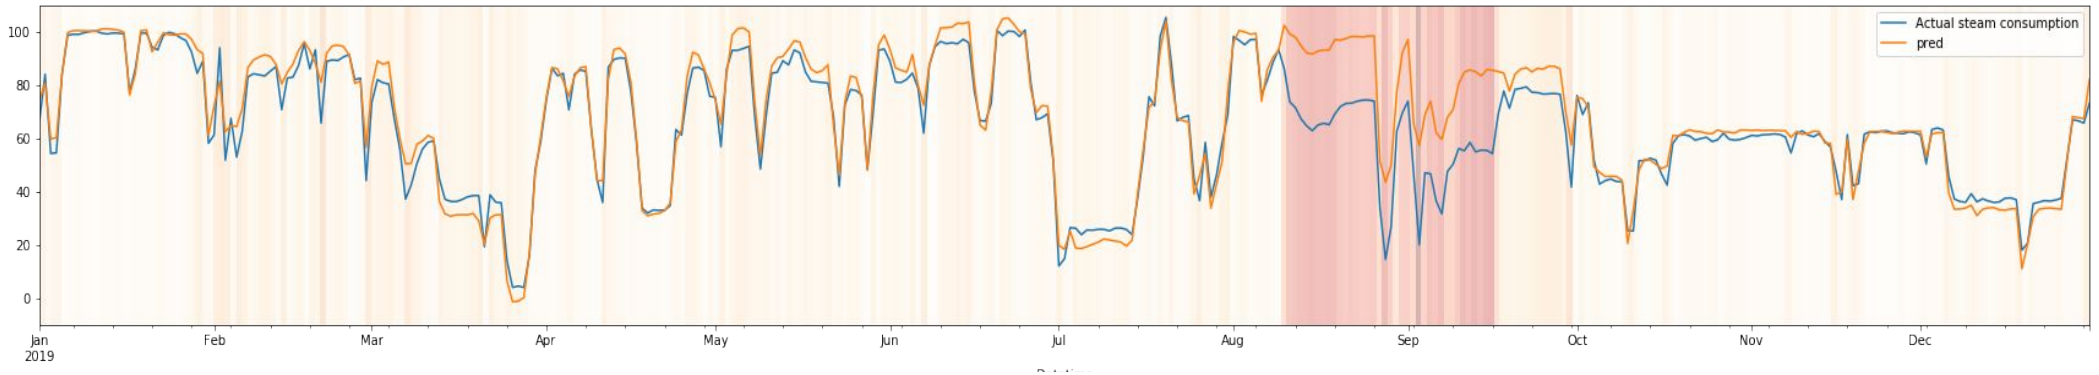

Total | Patent classification

Description

- There are a massive number of new patents published every year, the goal is to **automatically extract the most relevant / interesting** ones (subjective interestingness)
- There is a large databases of historic patterns with **labels** (patents that were looked into are labeled “interesting”
- Secondary goal is to **detect trending topics** (e.g. spot new technologies coming up)

Challenges and solutions

- **Feature engineering**
  - Patent metadata (e.g. publication date, IPC classification, …)
  - NLP: embeddings, tf-idf weighting, …
- Avoid **overfitting**
  - Simplify features
  - Further NLP filtering of noisy words
- General model performance
  - **Ensembling** of models with different feature sets

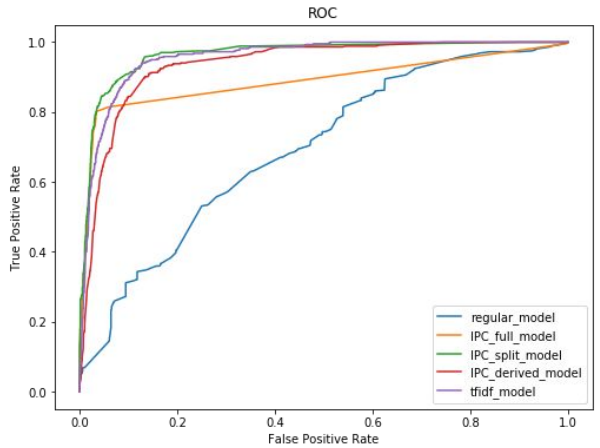

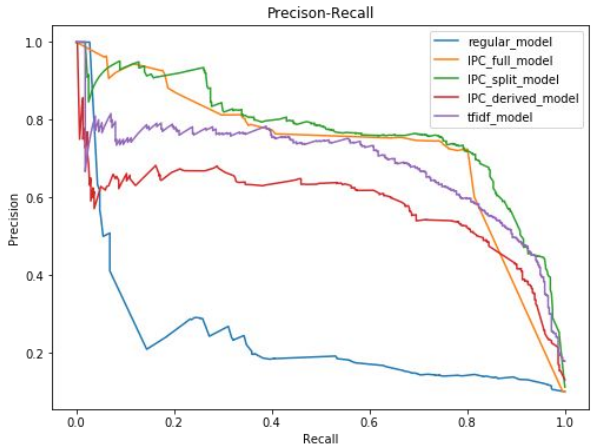

Co-libry | Similarity learning

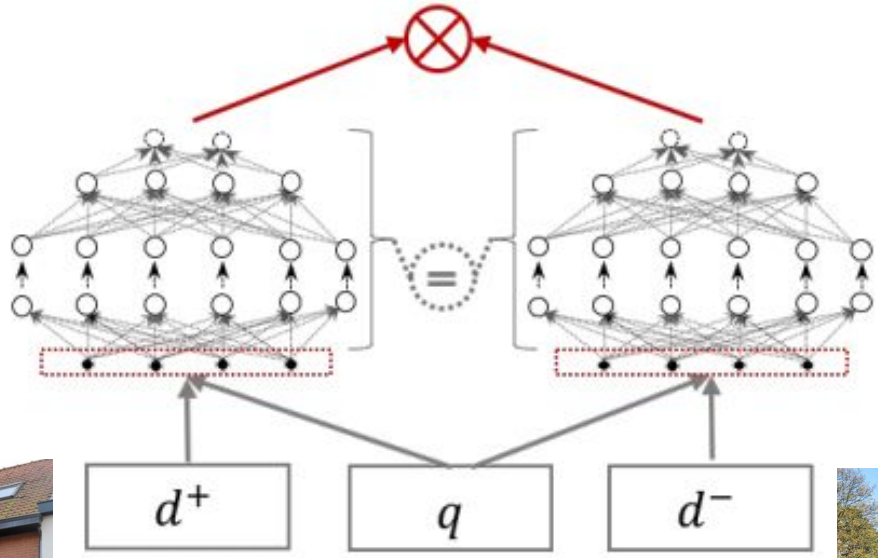

Positive example
listing viewed immediately after anchor

Anchor
listing viewed by user

Negative example
random listing

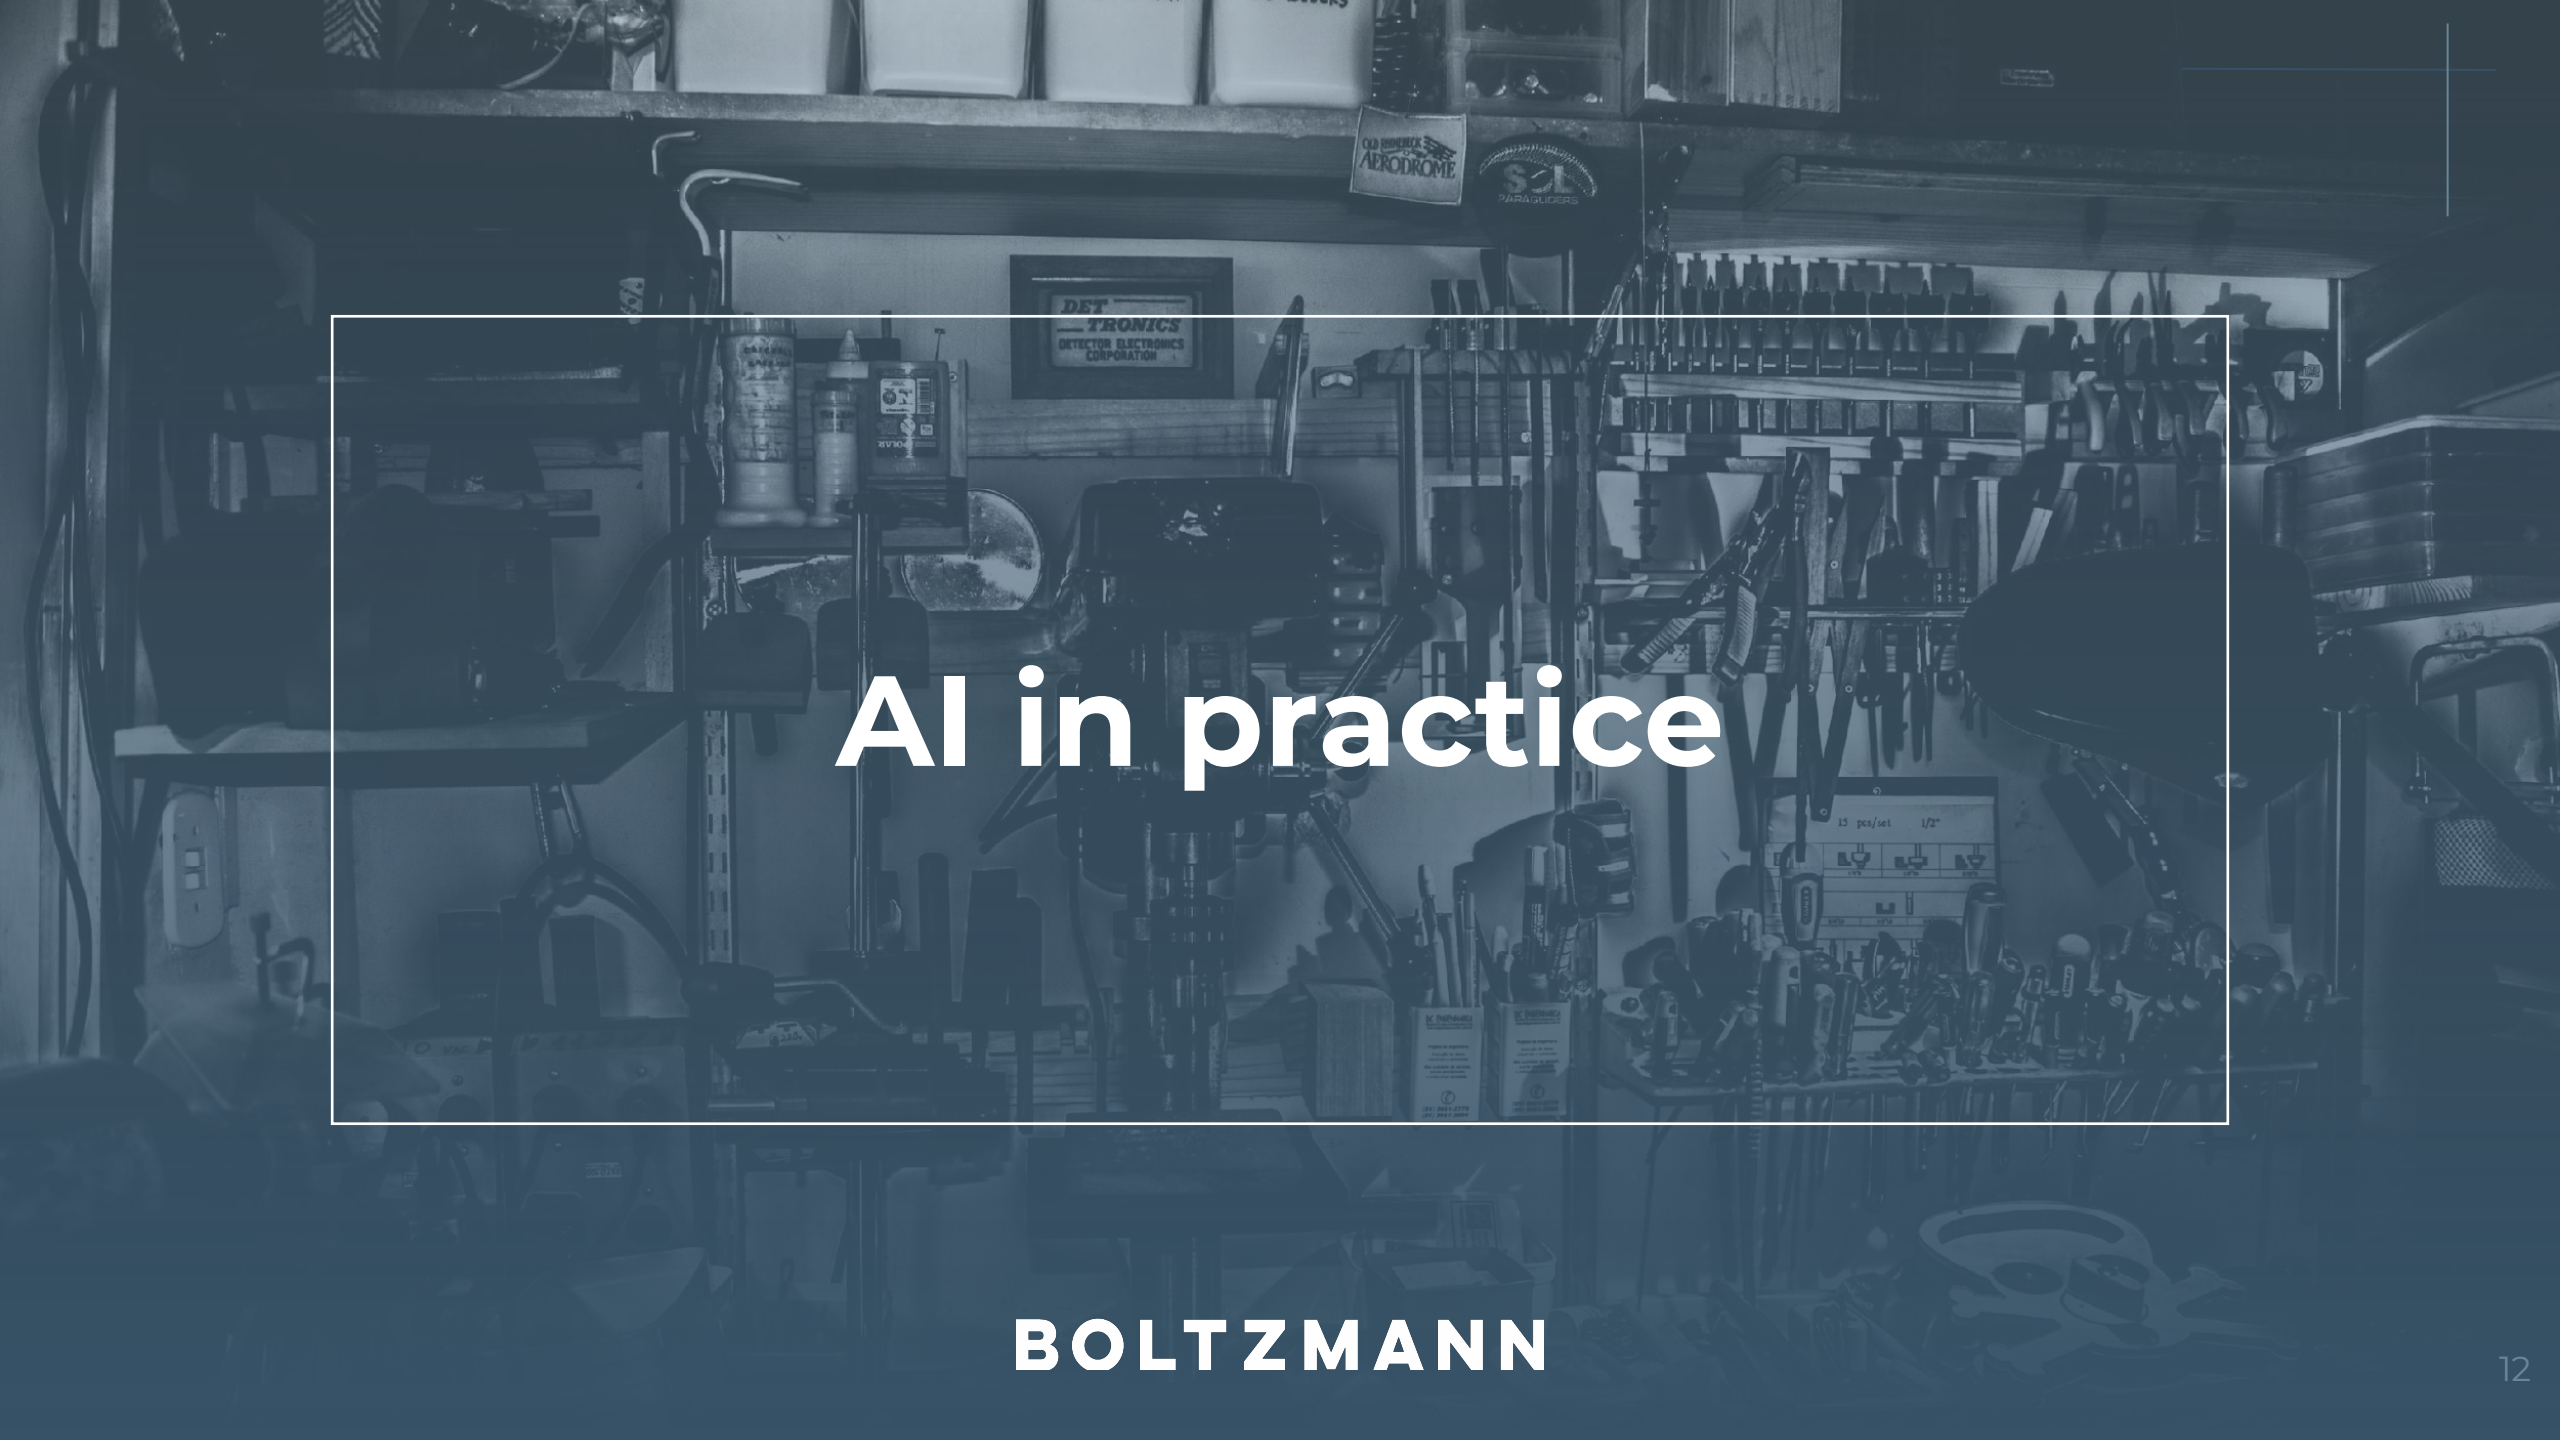

AI in practice

BOLTZMANN

Expectation	Reality
A clear goal A year to investigate apply all techniques available Client involvement	A vague idea / a vision, but no strategy Budget restrictions Skepticism
Given • target data • perfect, final version of the input data	Given • a vague idea • some data sets, in various places
Data in the cloud As many CPUs, GPUs and RAM as you need	Often servers on premise (security, data ownership) Limited to available hardware

Data sets: expectation/research vs. reality

	Feature 1	Feature 2
1		
2		
3		
4		
5		
...		
1e9		

	Feature 1	Feature 2	...	Feature 1000
1				
2				
3				

Bizarre quotes

- “We are neural network experts. This data set is too small, so we cannot use NN. However, I will talk about NN for the next 20 mins.”
- “Statistics is dead. [...] Machine learning finds correlations in the data.”
- “We have automated the data science process. With the press of a button, we fit 10k+ models and return the best result. No need for data scientists anymore.”
- “Machine learning is only about fitting.”
- “... So, then the model needs to refit every night, so it can make predictions for new data points.”
- “The machine learning model will become smarter every time it sees a new data point”
- “One of our guys knows a little bit about machine learning, because he installed TensorFlow.”
- “Now that we have all our documents digitized, we can press *the Artificial Intelligence button* and make our business A.I. driven.”
- “All structured data has been converted into value with machine learning. Now, businesses should focus on unstructured data only.”

LUDWIG

de A.I. Assistant voor Accountants

Shift from compliance to proactive advice

Traditional accounting

1. **Link** data
2. **Exact** calculation
3. **Manual** checks and review
4. Standardisation **per** client
5. **Expert** based advisory
6. Automate **booking & expense**
7. **(Near) real-time** reporting

AI powered accounting

1. **Interpret** data
2. **Complex** calculation with **uncertainty**
3. **Automated** checks and review
4. Standardisation **over** clients
5. **Automated** advisory generation
6. Automate **closing & post-accounting**
7. **Predictive** reporting

Assists mapping to an in-house charter of accounts

Ludwig Mapping

- Increased **consistency** across files
  → base for insights and review
- Automatic **reporting**
  → e.g. non-deductible expenses
- **Reduce time** consuming task
  → from hours to minutes

What Ludwig Mapping looks like **for the user**

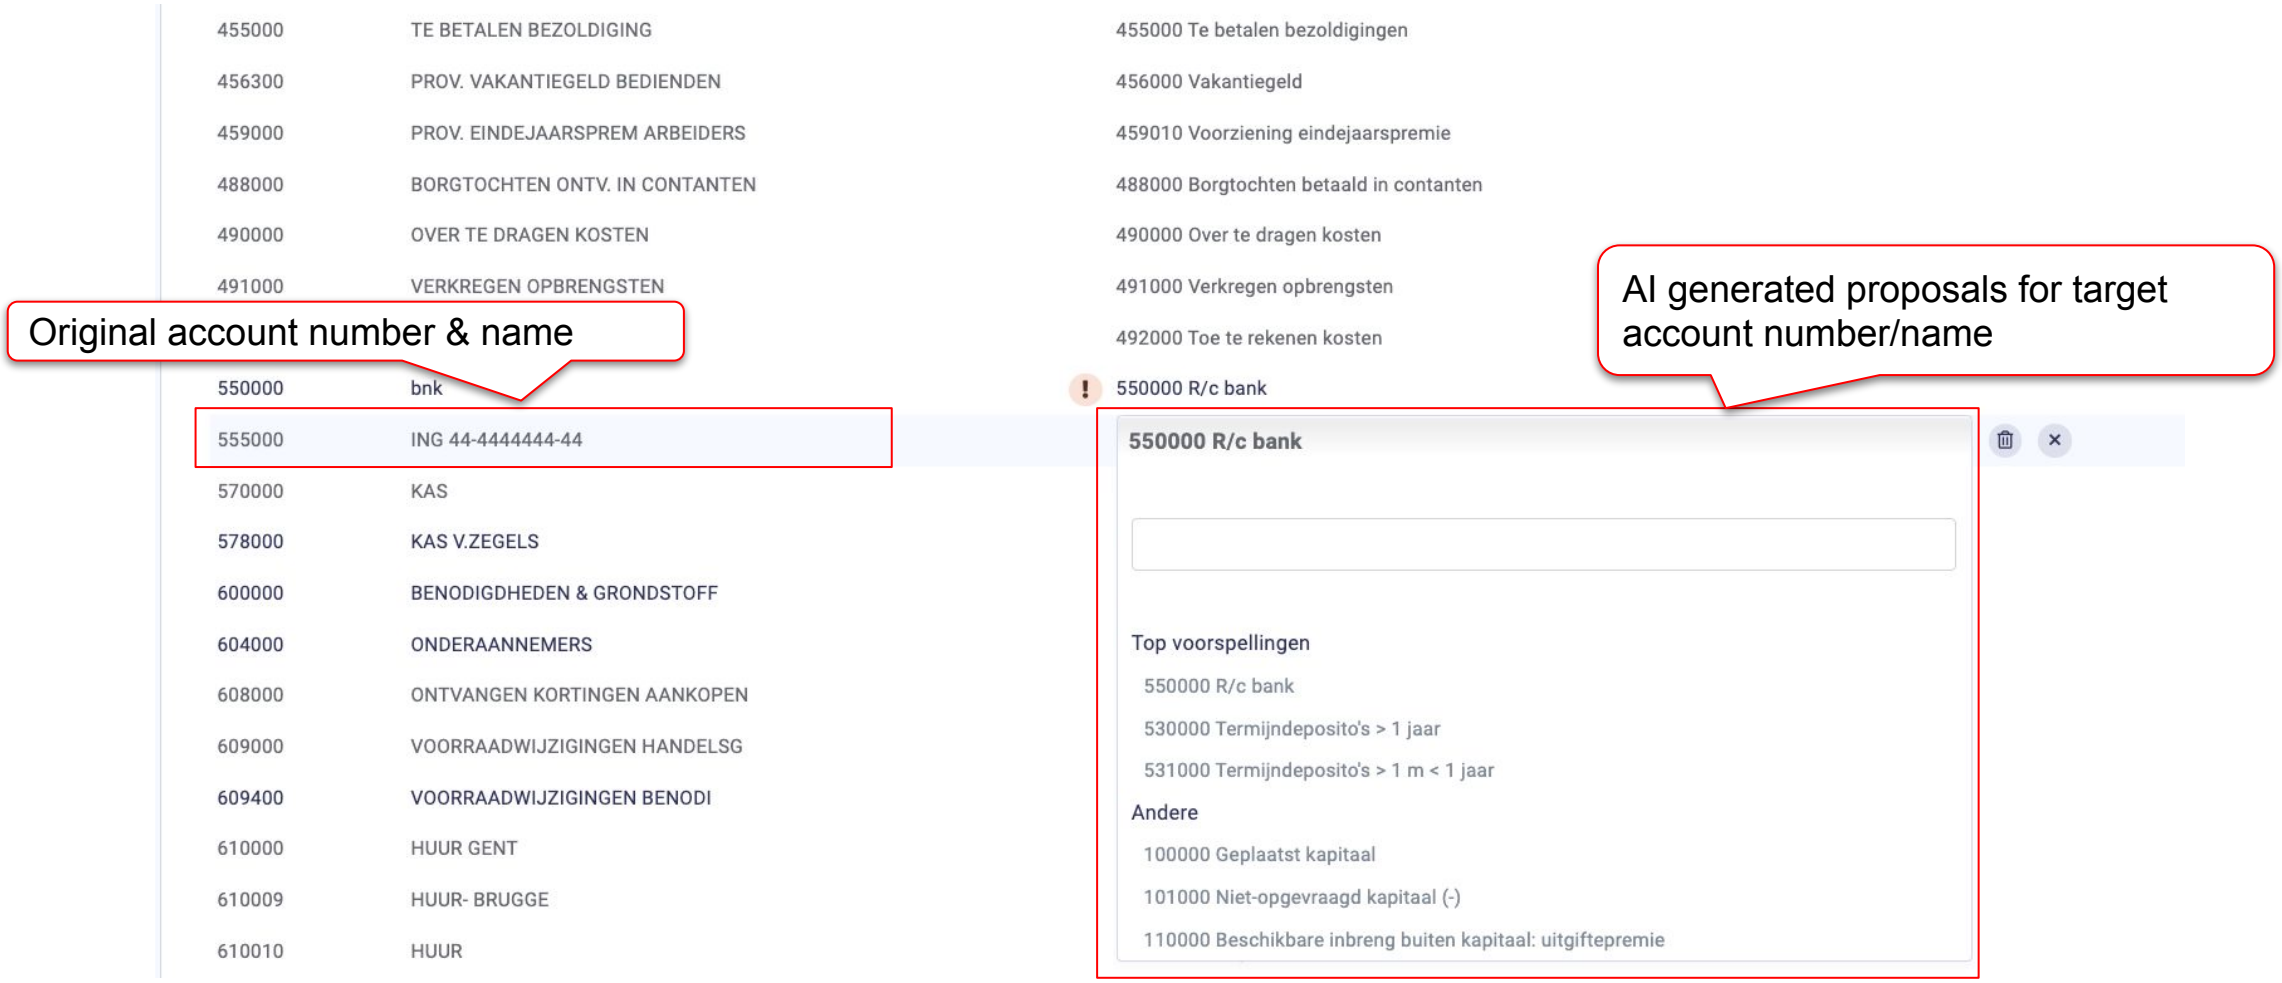

Ludwig | First a small intro to accounting

LUDWIG

Mapper | Review | Overview

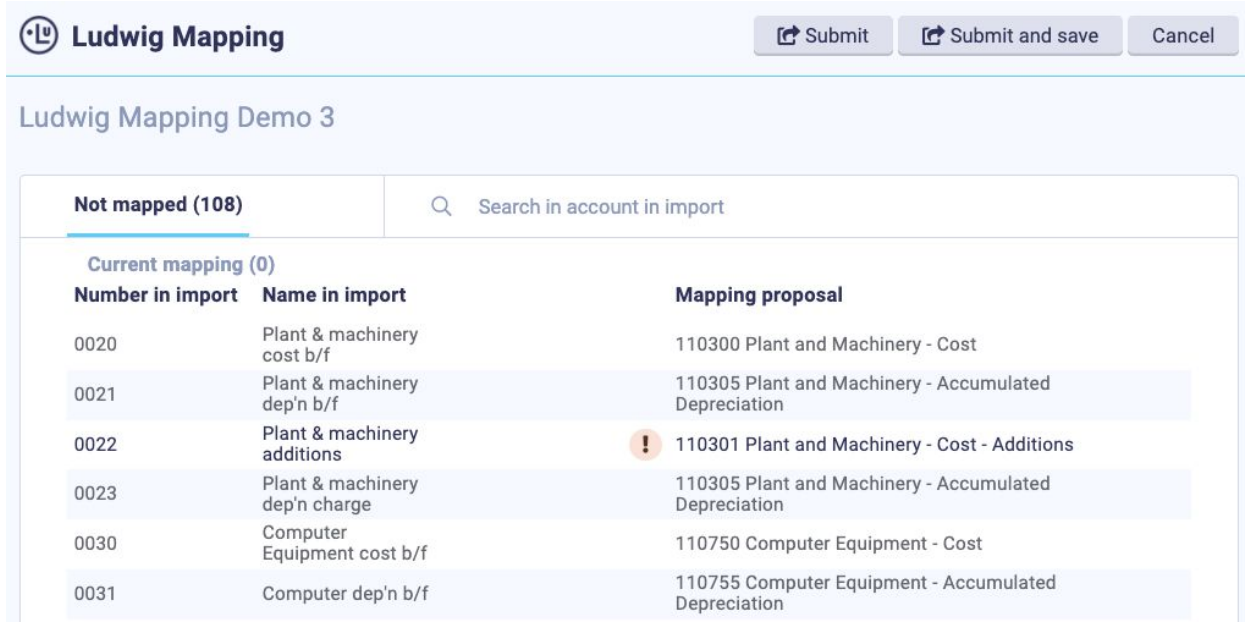

Ludwig | First a small intro to accounting

LUDWIG

Mapper | Review | Overview

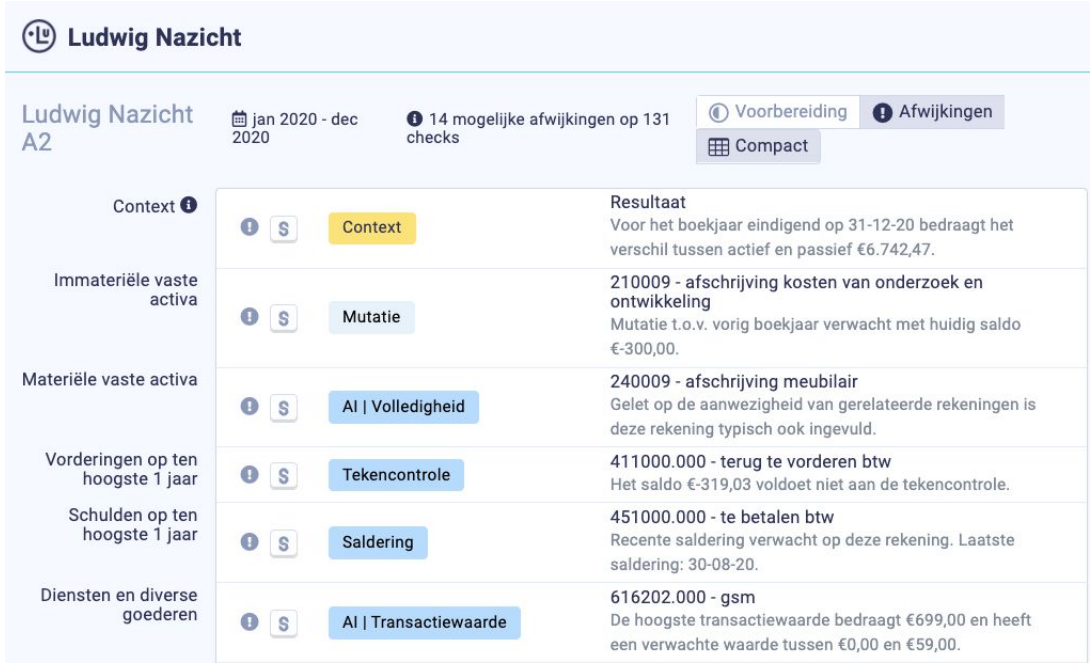

Ludwig | First a small intro to accounting

LUDWIG

Mapper

Review

Overview

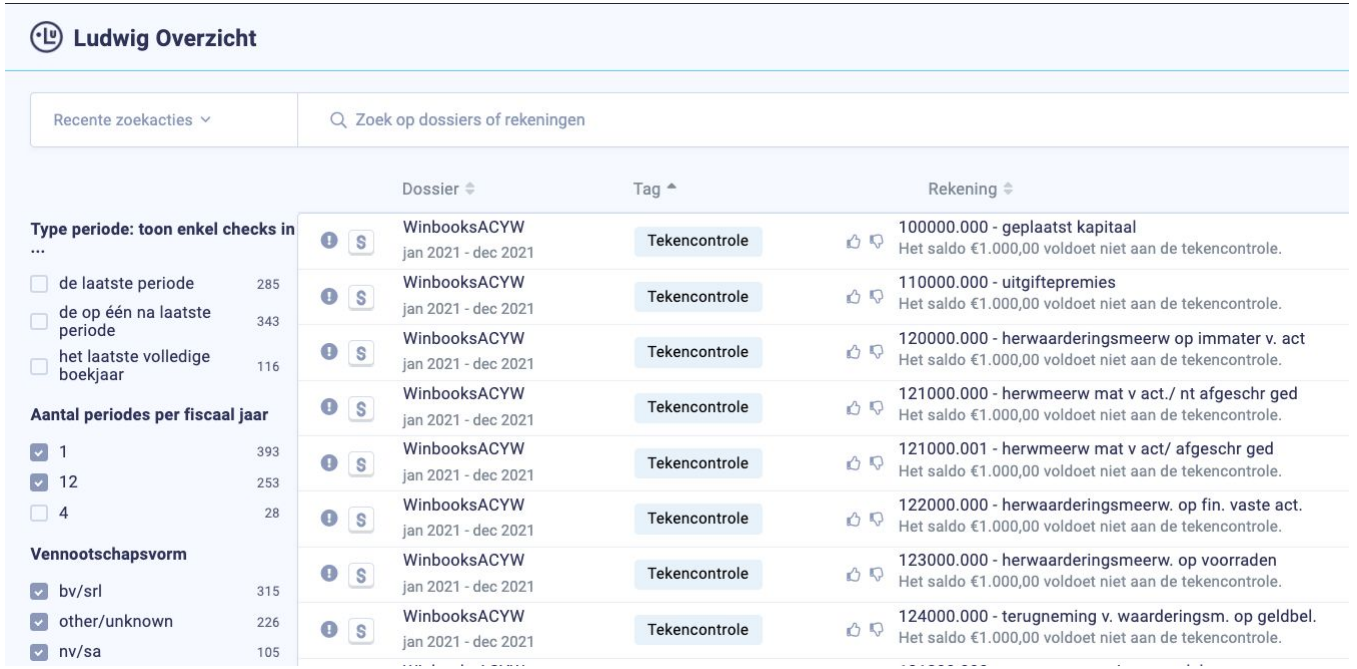

In short, Ludwig helps the accountant to …

LUDWIG
Mapping

… easily and quickly get consistent files and automatic reporting across all companies

LUDWIG
Review

… automatically detect mistakes and opportunities for advice in the financial data

Coming soon…

LUDWIG
Forecast

… track budgets and cash flow, optimise prepayments, etc by providing reliable forecasts

LUDWIG
Benchmark

… easily position their clients within a group of peers in order to extract insights

Thank you

0474 67 02 77
Thomas.Servotte
@boltzmann.be
Data Scientist

0488 24 43 31
Ken.Bastiaensen
@boltzmann.be
Co-founder & partner