WAN-IFRA Data Group Meeting

Causal Machine Learning @SZ

Felix Hagemeister – Data Scientist, Süddeutsche Zeitung Digitale Medien

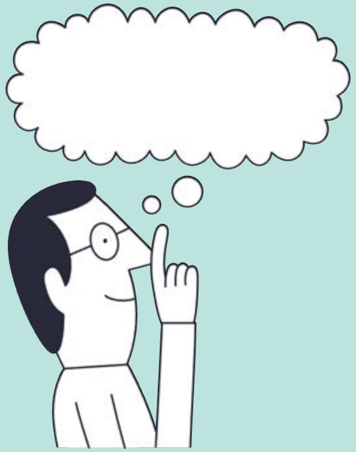

Süddeutsche Zeitung

Motivation

Why prediction does not yield actionable insights

These are two customers Herbert and Lara. You have built a machine learning model that predicts churn.

Your churn model predicts that **Herbert has a high churn probability.**

Your churn model predicts that **Lara has a low churn probability.**

Whom would you target with a retention marketing campaign? Why?

Motivation

Why prediction does not yield actionable insights

💡 **Key insight:** Customers identified as having high probability of churning are not necessarily the best targets for proactive churn measures (Ascarza 2018, Journal of Marketing Research).

⚙️ **Causal ML**

- Allows to find out for whom which measure has the highest leverage to prevent churn – independently from individual churn probabilities
- Extends AB-testing by personalizing A or B interventions to optimize uplift from interventions
- Gives actionable insights to effectively guide marketing measures.

Motivation

Why prediction does not yield actionable insights

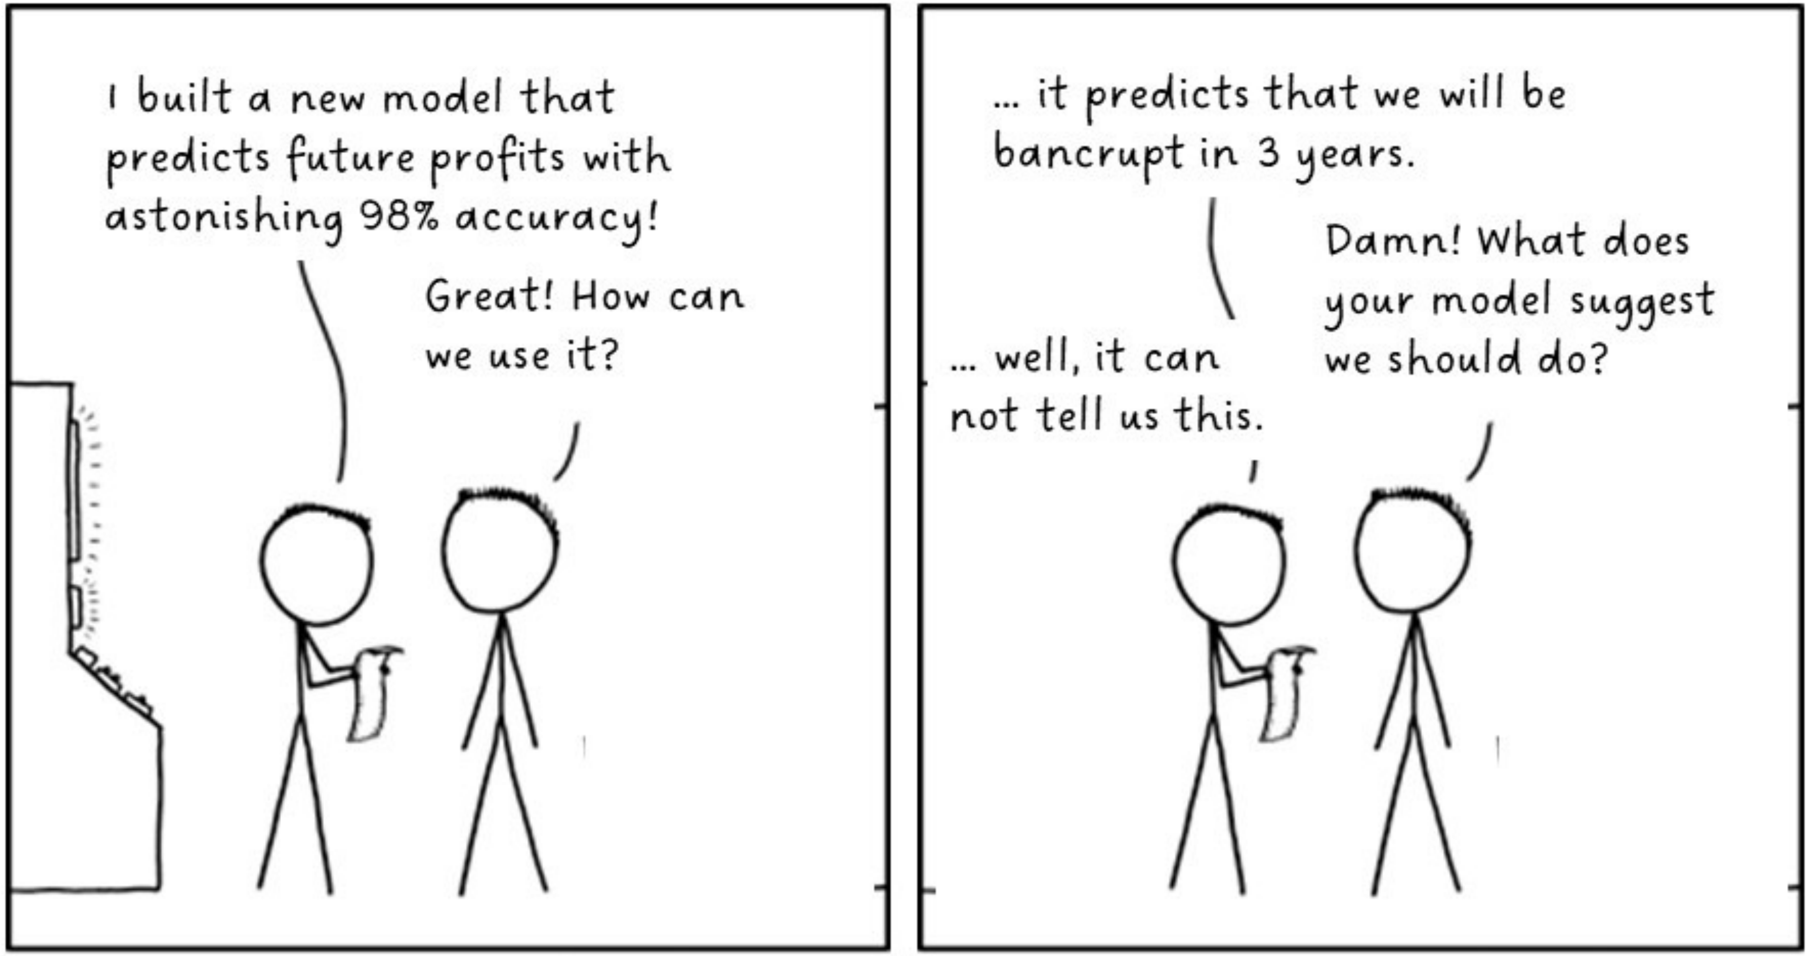

Süddeutsche Zeitung

Overview: Tools for Causal ML

How to combine prediction and causal inference

Main Packages:

CausalML (Uber)

EconML (Microsoft)

Goal: Estimate individual or group-based average treatment effect, aka "Uplift"

Algorithms:

- Tree-based algorithms: Causal Trees, Causal Forests
- Meta-learner algorithms: S-, T-, X-, R-, and DR-Learner
- Neural Nets: Causal Effect Variational Autoencoder and Dragonnet

Use Case I: Metering Test

How to personalize metering limit to increase conversions

AB-Test summer 2022: Should we increase the metering limit?

Goal: Increase subscription conversions

2,3M Users

Group A Metering = high

Group B Metering = low

Simple AB-Test Result: High metering **decreases** conversion probability by 0.01%.
(However, **higher** return rate and engagement)

Use Case I: Metering Test

Decision Tree

Traditional:

- build decision tree that models purchase propensity

Whom would you assign the low and the high metering limit? Why?

(example of a decision tree)

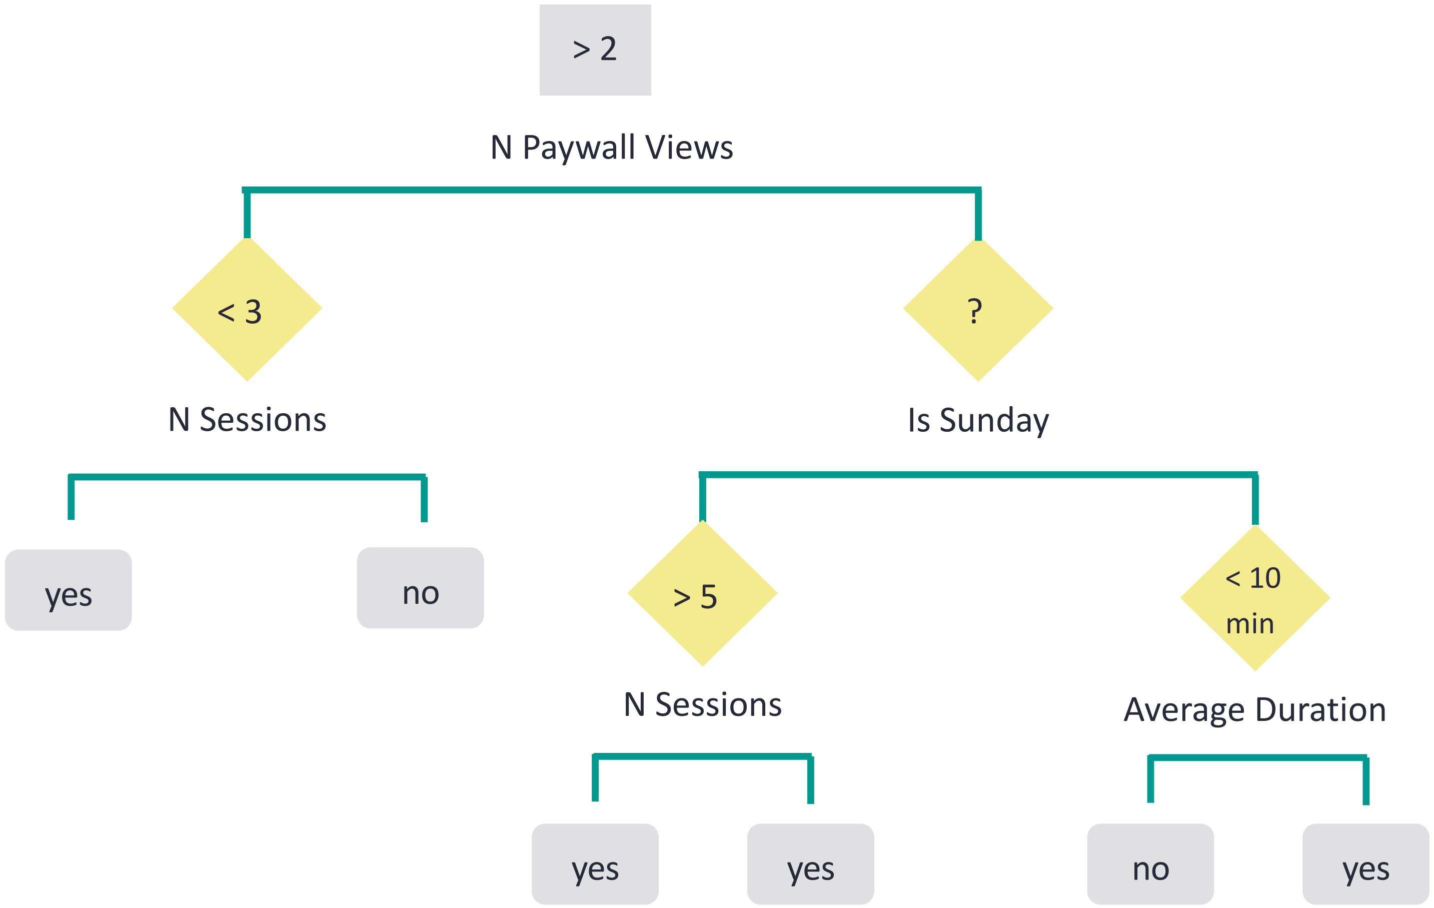

Use Case I: Metering Test

Causal Tree

A tree "grows" such that it builds distinct groups with different leverages of intervention ("uplift effects"):

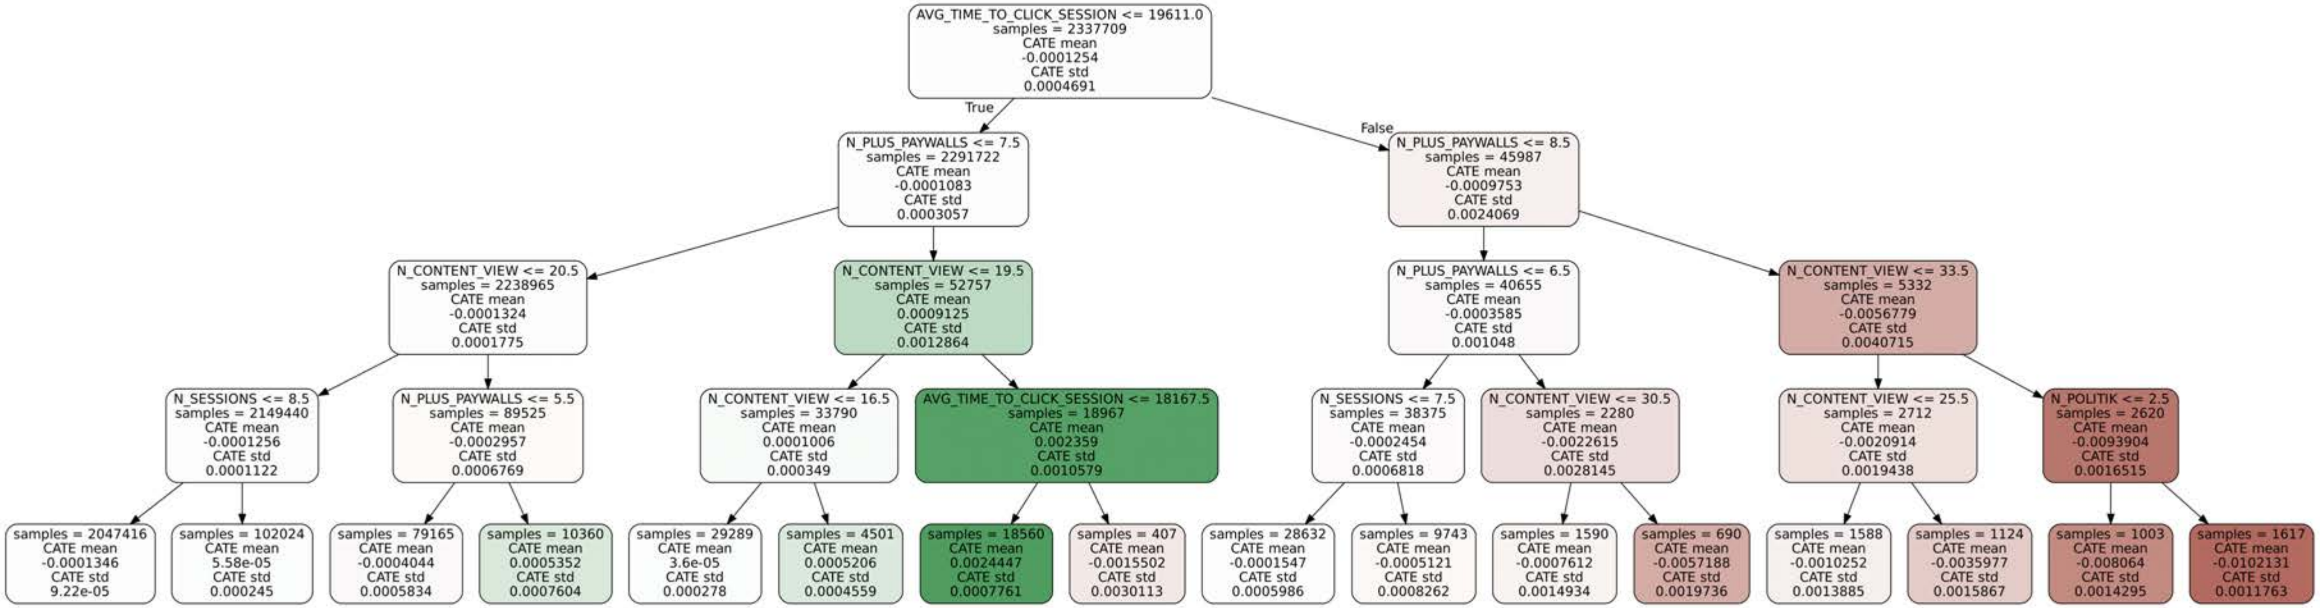

Some groups have positive uplift effects (green boxes)

Some groups have negative uplift effects (red boxes)

Use Case I: Metering Test

Uplift Effect Heterogeneity

Metering limits do not matter for conversion probabilities for most users (curve is 0)

Some users have higher conversion probabilities at the low metering limit (curve is negative)

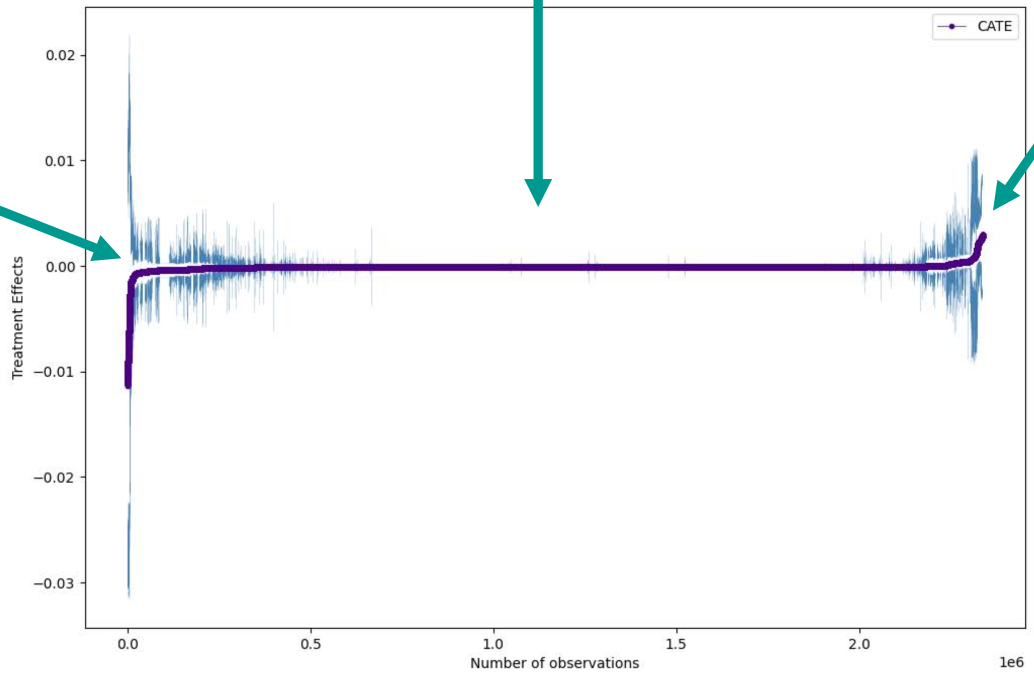

Some users have higher conversion probabilities at the high metering limit (curve is positive).

Use Case I: Metering Test

How to personalize metering limit to increase conversions

How many more conversions because of Causal Trees vs. ATE in Metering?

During the experiment (2 months, ~2.3M users) we could have reached...

Scenario	Result
Random split	*-150 conversions*
Metering = low (ATE)	>5000 conversions
Causal Trees	**+80 conversions**

Conversions **increase** by ~1.2%.

Use Case II: Onboarding

How to personalize onboarding e-mails to improve retention

AB-Test spring 2023: Should we change the onboarding e-mailing scheme?

Goal: Increase trial conversion rates (retention)

31k Users

Group A Old onboarding scheme

Group B New onboarding scheme

Simple AB-Test Result: New onboarding scheme **decreases** trial conversion rate by 1%.

Use Case II: Onboarding

Churn Prediction

Traditional:

- build an XGB model that predicts churn within trial subscription

Whom would you assign the old and the new onboarding scheme? Which cutoff would you pick?

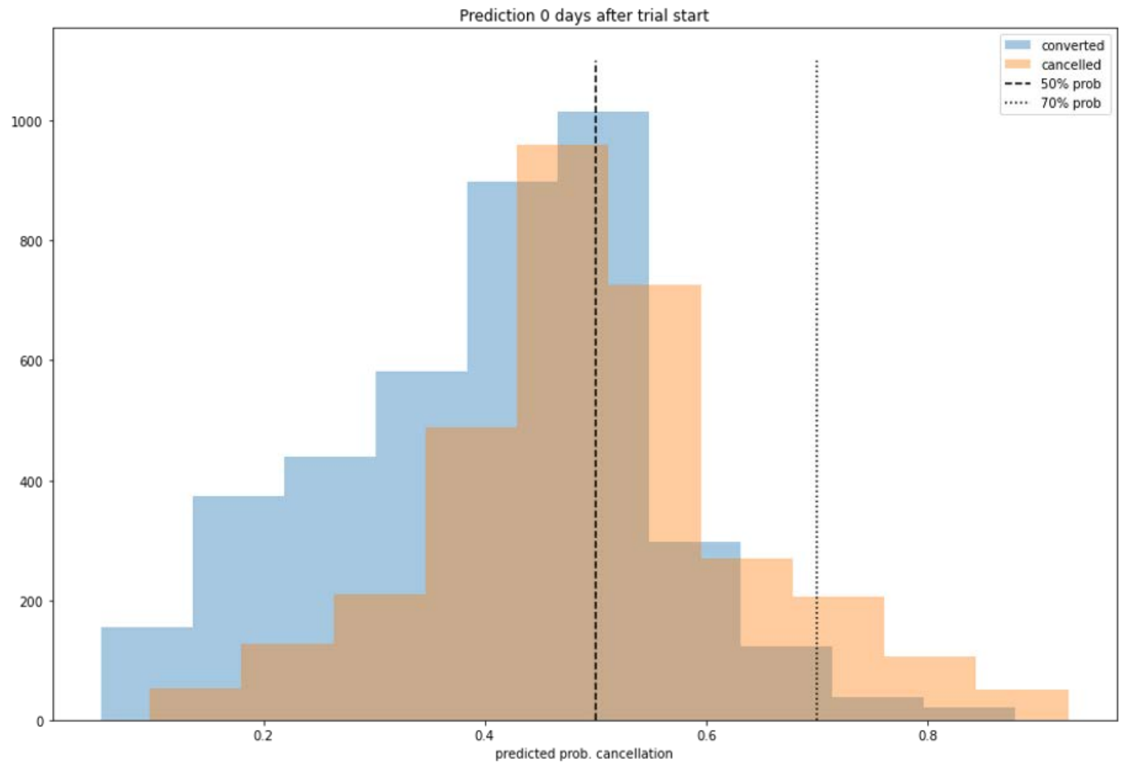

cutoff	precision	recall	f1-score	accuracy
0.5	59%	48%	53%	62%
0.7	80%	8%	15%	58%

Use Case II: Onboarding

How to personalize onboarding emails to improve retention

Causal ML delivers...

💡 **Actionable Insight:** Each onboarding should be used for people in selected target groups.

📈 **Impact:** Personalizing onboarding using causal trees **increases** retention rate by ~1.1 percentage points.

🔎 **Interpretability:** Feature combinations of selected target groups can be interpreted.

📏 **Scope:** ~ 11% of users should be assigned the new onboarding scheme.

Conclusions

Causal ML

- Simple prediction models are unable to draw actionable conclusions and it is hard to show actual impact of these models in production.
- Causal ML allows us to more effectively target our users according to high expected impact of campaigns or other interventions.
- Interventions can be of any kind – marketing campaigns, customer success measures, content recommendations, …