Setting optimal prices during a sales event steered by humans

Price optimization at Zalando

Adèle Gouttes

Berlin, Co@Work.zib
27.09.2024

Agenda

- Algorithmic Pricing: An Introduction
- Challenge 1: Human interactions
- Challenge 2: Data collection

00

Algorithmic Pricing

An Introduction

Algorithmic pricing: Introduction

Nike Sportswear

MANOA 17 UNISEX - High-top trainers

41,45 € VAT included

Originally: ~~54,95 €~~ -25%

★★★★☆ 1

Colour: **wheat/black**

Choose your size

Add to bag

1-2 working days

Algorithmic pricing: Goal

Goal:
Recommend the best discount for ~900K articles across 25 countries, such that business constraints are fulfilled.

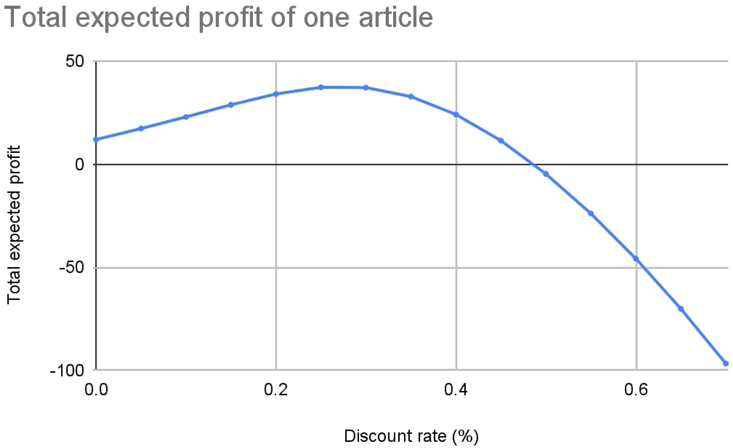

Algorithmic pricing: Goal

“Best discount”:
In order to maximize Zalando’s profit over the year.

“Business constraints”:
- Financial targets: revenue, profit, average discount rate
- Targets are set at country-level.

Algorithmic pricing: The recommendation tool

Basic setup:
- First-forecast-then-optimize
- Forecasts:
  - Demand and costs
  - For each article (~900K, 25 countries)
  - Future looking: “how much do we sell next week?”
  - As a function of discount: “how much do sell with 10%, 20%... discount?”
- Optimizer:
  - Selects the best discounts, in order to maximise profit
  - Leads to ~900K x 25 discount decisions

Algorithmic pricing: The ecosystem around the tool

The recommendation tool belongs to an ecosystem:

- **Pricing managers:**
  humans in charge of reaching financial targets
- **Data collection**:
  in order to train our ML models, we collect all kind of data.

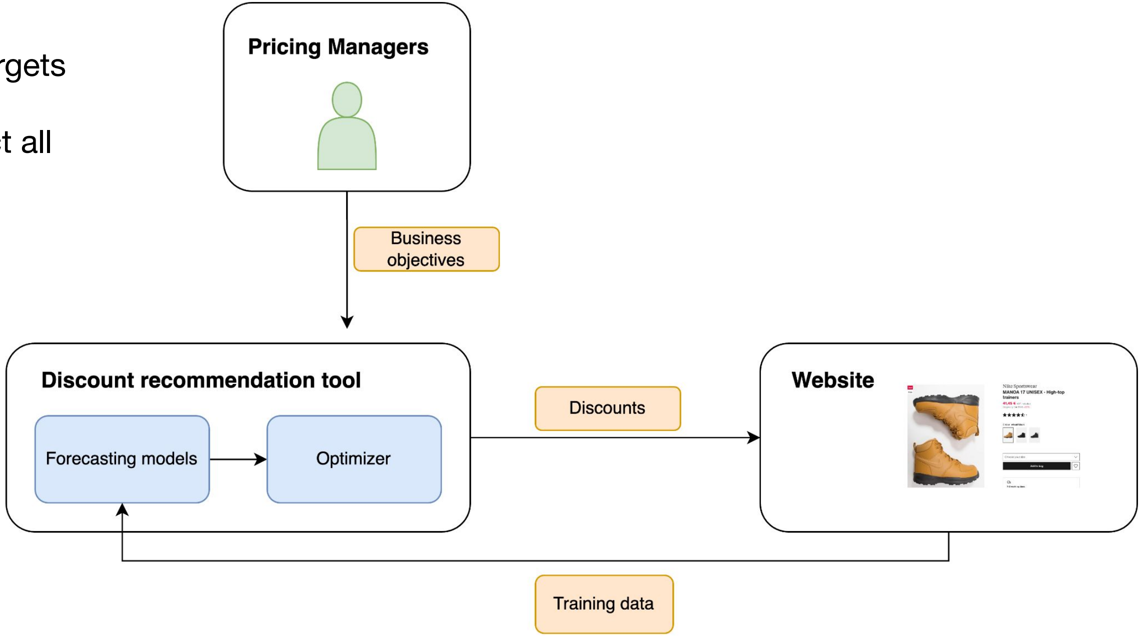

Algorithmic pricing: Challenges

- The recommendation tool is **not an isolated component.**
- Specific challenges come from its environment
- Additional requirements when designing the tool

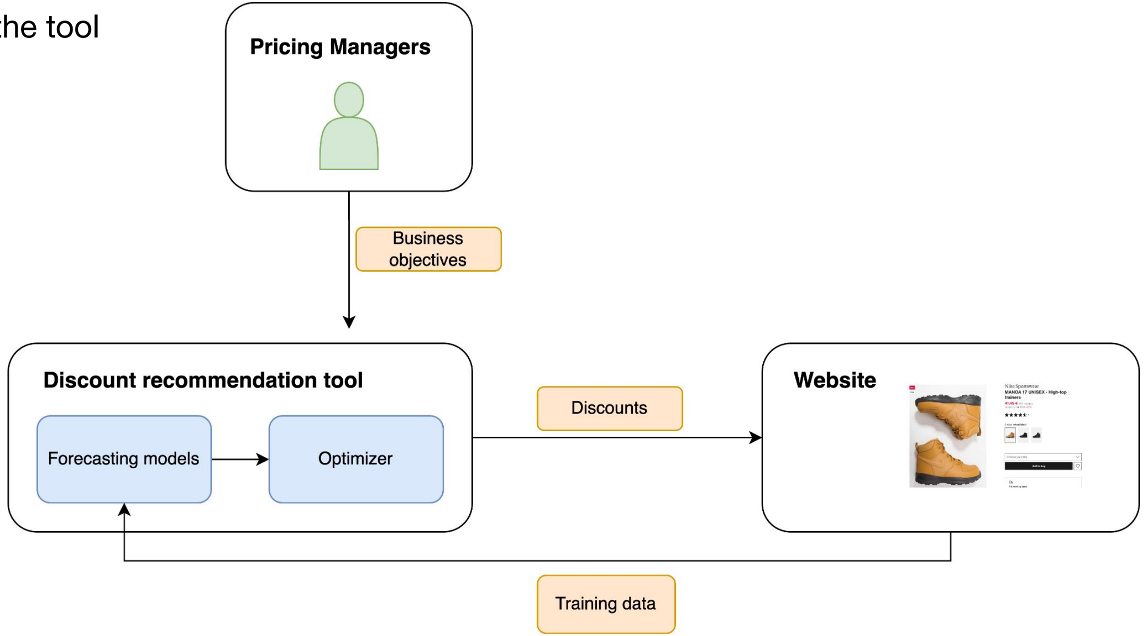

Algorithmic pricing: Challenges

Challenges:
1. **Human interaction**:
   Pricing managers need to “tell the tool what to do”
2. **Data collection**:
   Our past decisions bias our training data

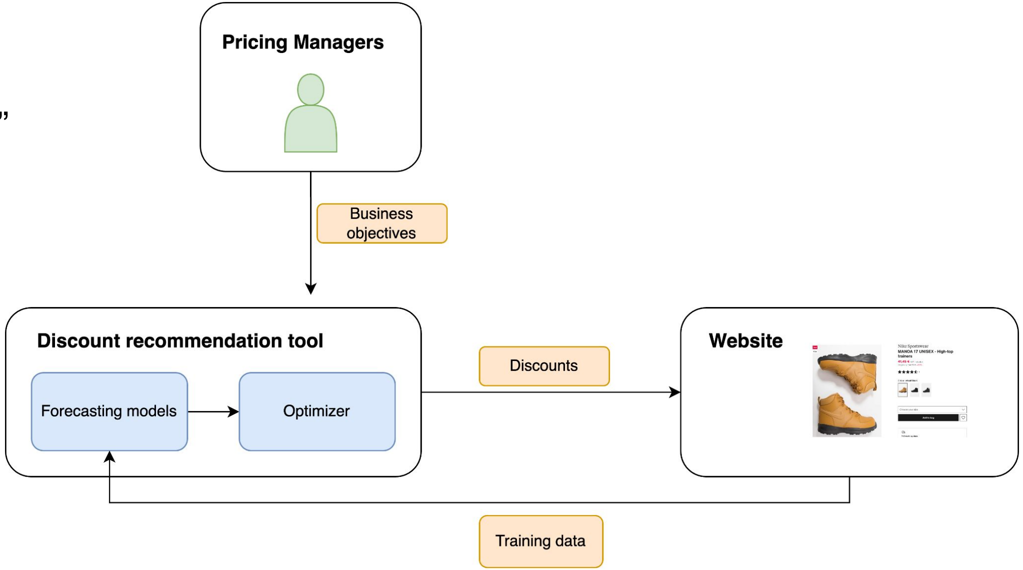

01

Challenge 1:
Human interactions

Human interactions: Motivations

1. **Responsibility**: the pricing manager is responsible for reaching the financial targets, not the tool

2. **Complex decision making in the company**: Pricing is just one leverage, the department needs to align with logistic, marketing, etc…

3. **Lack of data**: a competitor with an aggressive campaign, the euro-cup, etc…

A **good** tool requires **as little interaction as possible**, but some will always be necessary.

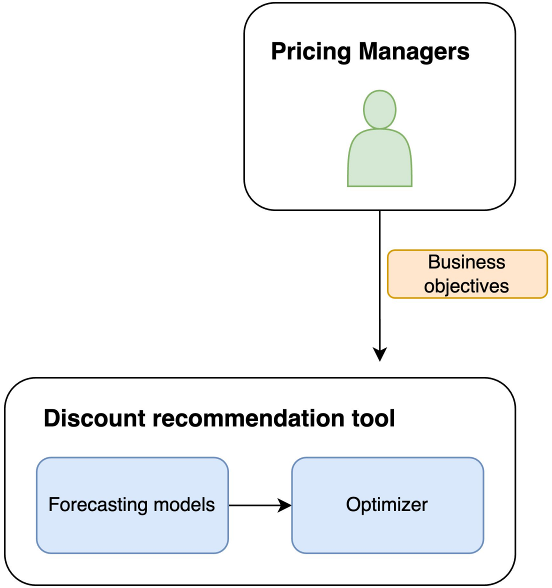

Human interactions: Design of the optimizer (the deprecated way)

Financial targets given as inputs to the optimizer:
Solve all together:

- Optimize article-level long term profit (~900K)
- While satisfying global constraints:
  - 1 target per country (25)

This is a complex problem:

- Computationally **expensive**
- The full matrix of targets might **not be reachable**:
  In case of conflict, arbitration is needed
- Imposes a **slow** cadence (bi-weekly):
  Sometimes, price managers need to react fast.

Human interactions: Design of the optimizer (the new way)

Simplify the optimization problem:

- Financial targets become the **output**
- Break the problem **into independent components:** 1 per article x country
- No constraint
- The objective is a trade-off between immediate revenue and long term profit

$$Max_d LongTermProfit(d) + \alpha \times Revenue(d)$$

Human interactions: Design of the optimzer (the new way)

The pricing managers **choose a unique input-parameter (alpha)**, to reach the country financial target.

$$Max_d LongTermProfit(d) + \alpha \times Revenue(d)$$

Challenges:

- Alpha has **only a theoretical meaning** (related to marginal return on investment)
- In practice, alpha cannot be observed.
- This means that the pricing managers need **helpers** to take decisions.

A map between alpha and the predicted financial target

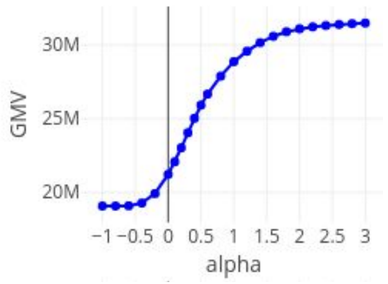

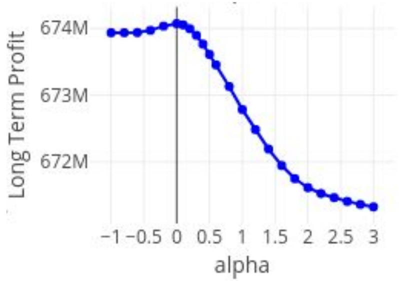

Human interactions: Learnings

Learnings:

- The attempt to satisfy all business constraints can lead to unnecessary complex systems.
- Always challenge a constraint/assumption that adds complexity: What is the impact?

02

Challenge 2: Data collection

Data collection: Training a ML model

The feedback loop:

- Train the forecasting models on Zalando sales
- Recommend the best discounts according to the models
- Upload the discounts on the website
- Collect data

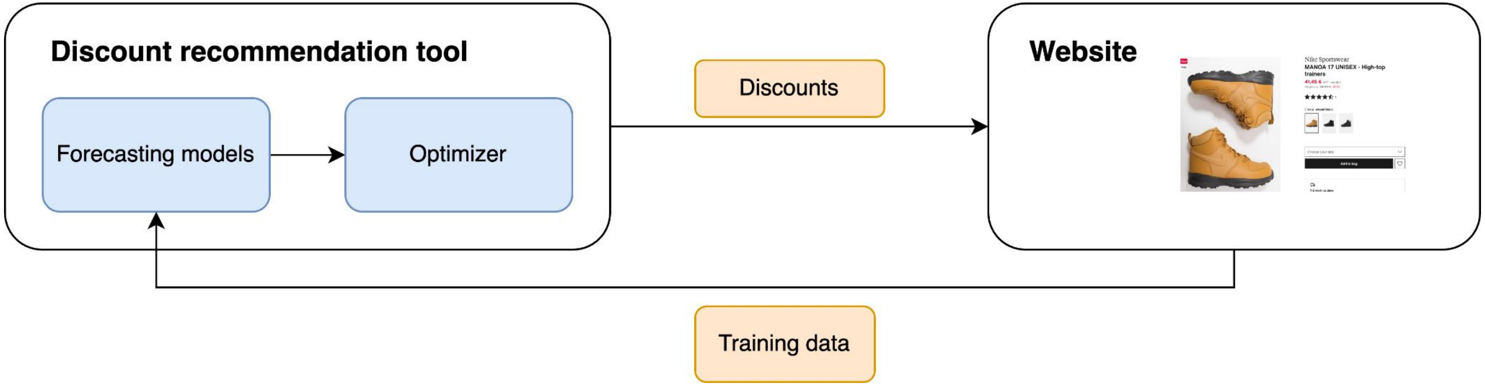

Data collection: Training a ML model

Challenge:

- ML models are very good at **extrapolating.** They learn correlations.
- But in practice we never observe some situations:
  - Best-sellers are never highly discounted
  - Low-sellers receive very high discounts

→ **So how is the model extrapolating exactly??**

- Our training data is **biased**
- Our **evaluation data is biased too**: we are blind

Data collection: The danger of bad extrapolation

An opportunity loss:
We might have earned more profit from this article, had we chosen a different price.

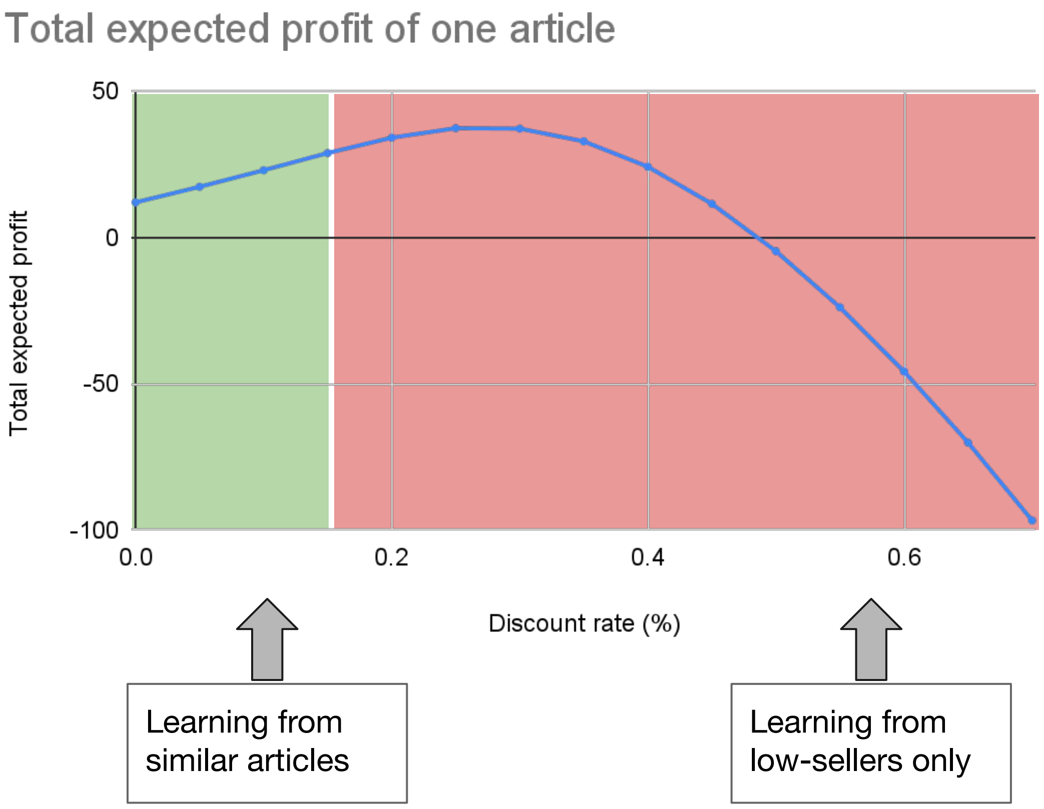

Data collection: Causation is not correlation

The ML model should predict the **causal effect:**
“How many euros of profit do I earn by choosing discount X?”

Independently of all confounders (article type, time of the year, etc…)

Data collection: How to mitigate the bias problem?

How to mitigate the problem?

- **Forbid** some decisions: “the best articles shall not get more than 10% discount”.
- **AB tests**:
  some decisions should be carefully tested before deployment.
- **Causal ML**:
  Coming from economics, how to build “what if” scenarios.

And many others: stochastic optimization, simulation…

Data collection: Learnings

Learnings:

- In a forecast-then-optimise set up, we might create a biased feedback loop
- Causation is not correlation: our ML model should be causal

03

Conclusion

Conclusion

- Zalando optimize prices using a forecast-then-optimise set up
- A recommendation tool cannot be designed in isolation
- The environment creates challenges, and naive solutions generate losses
- Two examples: human interaction, data collection

But we have many more challenges :)

References

Kunz, M., S. Birr, M. Raslan, L. Ma, Z. Li, A. Gouttes, M. Koren, et al. 2023. “Deep Learning based Forecasting: a case study from the online fashion industry.” *In Forecasting with Artificial Intelligence: Theory and Applications* (Switzerland), 2023.

Streeck, R., T. Gellert, A. Schmitt, A. Dipkaya, V. Fux, T. Januschowski, and T. Berthold. 2024. “Tricks from the Trade for Large-Scale Markdown Pricing: Heuristic Cut Generation for Lagrangian Decomposition.”

Schultz, D., J. Stephan, J. Sieber, T. Yeh, M. Kunz, P. Doupe, T. Januschowski. 2024 “Causal forecasting for pricing”