

Marketing: Selected Doctoral Theses

“Product Returns Management in Online Retail”

Author: Marat Ibragimov (2023)

Committee: John R. Hauser (chair), Duncan Simester, Artem Timoshenko (Kellogg School of Management, Northwestern University)

Abstract:

In Chapter 1, I and coauthors study the problem of predicting the product return rate using the products’ visual information. In online channels, products are returned at high rates. Shipping, processing, and refurbishing are so costly that a retailer’s profit is extremely sensitive to return rates. Using a large dataset from a European apparel retailer, we observe that return rates for fashion items bought online range from 13% to 96%, with an average of 53% – many items are not profitable. Because fashion seasons are over before sufficient data on return rates are observed, retailers need to anticipate each item’s return rate prior to launch. We use product images and traditional measures available prelaunch to predict individual item return rates and decide whether to include the item in the retailer’s assortment. We complement machine-based prediction with automatically extracted image-based interpretable features. Insights suggest how to select and design fashion items that are less likely to be returned. Our illustrative machine-learning models predict well and provide face-valid interpretations – the focal retailer can improve profit by 8.3% and identify items with features less likely to be returned. We demonstrate that other machine-learning models do almost as well, reinforcing the value of using prelaunch images to manage returns.

In Chapter 2, I consider customer search and product returns on the individual level. Previous research has focused on linking customers’ purchase and return decisions. However, online retailers have access to the information which precedes the purchase decision – customer search. I demonstrate that customer search information provides important insights about product returns. Using data from a large European apparel retailer, I propose and estimate a joint model of customer search, purchase, and return decisions. I then provide theory and data indicating that using search filters, viewing multiple colors of a product, spending more time, and purchasing the last item searched are negatively associated with the probability of a return. Finally, I use the proposed model to optimize the product display order on the retailer’s website.

Chapter 3 extends and reinforces the results obtained from previous Chapters. In the paper, I study the assortment planning problem in presence of frequent product returns. I develop a deep-learning model of customer search, purchase, and return. The model is based on a transformer framework and allows the recovery of important relations in the data. I use the estimated model to demonstrate that retailers could identify successful and unsuccessful products and modify the assortment. The modified assortment would increase the retailer’s sales and at the same time decrease returns. Lastly, I provide qualitative insights on which products are most likely to be unsuccessful in online retail.

“Augmented Machine Learning and Optimization for Marketing”

Author: Yuting Zhu (2022)

Committee: Juanjuan Zhang (co-chair), Duncan Simester, Tony T. Ke (The Chinese University of Hong Kong)

Abstract:

This dissertation consists of three essays exploring how to augment machine learning and optimization methods for marketing management. The first essay considers an augmentation of deep-learning-based recommender system for sales force management. Helping new salespeople succeed is critical for many organizations. We develop a deep-learning-based recommender system to help new salespeople recognize suitable customers, leveraging historical sales records of experienced salespeople. One challenge is how to learn from experienced salespeople’s own failures, which are prevalent but often do not show up in sales records. We develop a parsimonious model to capture these “missing by choice” sales records and incorporate the model into a neural network to form an augmented, deep-learning based recommender system. We validate our method using sales force transaction data from a large insurance company. Our method

outperforms common benchmarks in prediction accuracy and recommendation quality, while being simple, interpretable, and flexible. We demonstrate the value of our method in improving sales force productivity.

The second essay explores an augmentation of large-scale linear programming optimization method for targeting with constraints. Personalization, which aims to target different marketing actions to different customers, has attracted broad attention in both academia and industry. While most research has focused on training personalization policies without constraints, in practice, many firms face constraints when implementing these policies. For example, firms may face volume constraints on the maximum or minimum number of actions they can take, or on the minimum acceptable outcomes for different customer segments. They may also face fairness constraints that require similar actions with different groups of customers. These constraints can introduce difficult optimization challenges, particularly when the firm intends to implement personalization policies at scale. Traditional optimization methods face challenges solving large-scale problems that contain either many customers or many constraints. We show how recent advances in linear programming can be adapted to the personalization of marketing actions. We provide a new theoretical guarantee comparing how the proposed method scales compared to state-of-the-art benchmarks (primal simplex, dual simplex and barrier methods). We also extend existing guarantees on optimality and computation speed, by adapting them to accommodate the characteristics of personalization problems. We implement the proposed method, and compare it with these benchmark methods on feasibility, computation speed, and profit. We conclude that, volume and similarity (fairness) constraints should not prevent firms from optimizing and implementing personalization policies at scale.

The third essay studies collective search in an organization. In this paper, we build a two-member two-period model to show that when a group of people with different preferences conduct search and make a decision together, they can benefit from making a commitment on the number of products to search ex ante when the search cost is very small or relatively large. The underlying mechanism is that, because of the preference divergence between group members, they tend to search fewer products and thus have lower expected utility in group search than in single agent search, and making a commitment on the number of products to search can help mitigate the preference divergence problem in group search. If consumers can observe product prices before search and the firm sets product prices endogenously, the firm can benefit from letting consumers commit to the number of products to search ex ante if consumers search as a group and their search cost is small. We also consider several extensions to show the robustness and boundary conditions of our findings.

“Essays on MarTech: Learning to Design, Deliver, and Diffuse Interventions”

Author: Jeremy Yang (2021)

Committee: Juanjuan Zhang (co-chair), Sinan Aral (co-chair), Dean Eckles

Abstract:

Chapter one develops an algorithm to predict the causal effect of influencer video advertising on product sales. A summary statistic, motion-score, or m-score, is proposed to capture the extent to which a product is advertised in the most engaging parts of a video. Pixel-level product placement is located with an object detection algorithm and pixel-level engagement is estimated as a saliency map by fine-tuning a deep 3D convolutional neural network on video-level engagement data. M-score is then defined as pixel-level engagement-weighted advertising intensity of a video. The algorithm is constructed and evaluated with influencer video ads on TikTok. Causal effects of video ads on product sales are identified by exploiting variation in video posting time. Videos of higher m-score indeed lift more sales. This effect is sizable, robust, and more pronounced among impulsive, hedonic, or inexpensive products. The mechanism can be partially traced to influencers' incentives to promote themselves rather than the product. How various stakeholders in entertainment commerce can use m-score in a scalable way to optimize content, align incentives, and improve efficiency are discussed.

Chapter two proposes a method to optimize a targeting policy that maximizes an outcome observed only in the long term. Traditionally, this typically requires delaying decisions until the outcome is observed or relying on simple short-term proxies for the long-term outcome. The method builds on the statistical surrogacy

and off-policy learning literature to first impute the missing long-term outcomes and then approximate the optimal targeting policy on the imputed outcomes via a doubly robust approach. It is applied in large-scale proactive churn management experiments at The Boston Globe by targeting optimal discounts to its digital subscribers to maximize their long-term revenue. It is shown that conditions for the validity of average treatment effect estimation with imputed outcomes are also sufficient for valid policy evaluation and optimization; furthermore, these conditions can be somewhat relaxed for policy optimization. The method is also validated empirically by comparing it with a policy learned on the ground truth long-term outcomes, they are shown to be statistically indistinguishable. It also outperforms a policy learned on short-term proxies for the long-term outcome.

Chapter three investigates how network embeddings can be applied to the study of product diffusion. Three sets of results are documented using a combination of real and simulated datasets: First, node embeddings can predict adoption decisions and timing better than standard centrality-based summary statistics. Second, node embeddings as control variables reduce the bias in the estimation of peer effect, especially when tie formation depends on unobservables. Third, graph embeddings based on the diffusion process as a whole reveal meaningful similarities between different diffusion processes such as simple vs. complex contagion.

“Machine Learning Methods for Targeting and New Product Development”

Author: Artem Timoshenko (2019)

Committee: John Hauser (co-chair), Duncan Simester (co-chair), Dean Eckles

ABSTRACT:

The dissertation consists of four essays on the applications of machine learning methods to targeting and product development. The first essay addresses the problem of identifying customer needs from user generated content. Traditionally, market research relies on interviews and focus groups to identify customer needs. User-generated content (UGC), such as online reviews, social media, and call-center data, provides an opportunity to identify customer needs more efficiently. Established methods are not well-suited for large UGC datasets because much of the content is uninformative or repetitive. We propose a machine learning approach for identifying customer needs from UGC and evaluate the method using a custom dataset matching Amazon reviews for oral care products to the customer needs identified from the interviews by professional analysts.

The second essay addresses the problem of training robust targeting policies. Targeting policies are used in marketing to match different firm actions to different customers. For example, retailers want to send different promotions to different customers, real estate agents want to show different homes, and car dealers want to propose different prices. We conduct two large-scale field experiments to evaluate seven methods widely used to design targeting policies. The findings compare the performance of the targeting methods and demonstrate how well the methods address common data challenges. The challenges we study are covariate shift, concept shift, information loss through aggregation, and imbalanced data. We show that more complex methods perform better when the training data is ideal, but they also deteriorate faster in the presence of the challenges that affect the quality of the training data.

The third essay addresses the question of efficiently evaluating targeting policies. Firms typically compare the performance of different targeting policies by implementing the champion versus challenger experimental design. These experiments randomly assign customers to receive marketing actions recommended by either the existing (champion) policy or the new (challenger) policy, and then compare the aggregate outcomes. We discuss an alternative experimental design and propose an estimation approach to improve the evaluation of targeting policies. The alternative experimental design randomly assigns customers to marketing actions. This allows evaluation of any targeting policy without requiring an additional experiment, including policies designed after the experiment is implemented. The proposed estimation approach identifies customers for whom different policies recommend the same action and recognizes that for these customers there is no difference in performance. This allows for a more precise comparison of the policies. We illustrate the advantages of the alternative experimental design and the proposed estimation approach using data from an actual field experiment.

