

A Multi-class Ride-hailing Service Subsidy System Utilizing Deep Causal Networks

Zhe Yu, Chi Xia, Shaosheng Cao (✉), and Lin Zhou

DiDi Chuxing

yzae2623@gmail.com, {xiachi, shelsoncao, realzhoulin}@didiglobal.com

Abstract. In the ride-hailing industry, subsidies are predominantly employed to incentivize consumers to place more orders, thereby fostering market growth. Causal inference techniques are employed to estimate the consumer elasticity with different subsidy levels. However, the presence of confounding effects poses challenges in achieving an unbiased estimate of the uplift effect. We introduce a consumer subsidizing system to capture relationships between subsidy propensity and the treatment effect, which proves effective, maintaining a lightweight online environment.

Keywords: Uplift modeling · Deep causal networks · Ride-hailing

1 Introduction

For online ride-hailing services, a pivotal marketplace strategy involves dynamically balancing demand and supply through the judicious use of subsidies for both consumers and drivers [2–8]. It is common practice to distribute subsidies to consumers upon receiving queries, which include the origin, destination, desired departure time, and other pertinent information. Subsequently, consumers make decisions on whether to proceed with the payment process, taking into account the subsidy offered to them. The overarching marketplace strategy aims to leverage a limited financial budget to expand the market scale.

Causal inference provides a fundamental approach to modeling consumer subsidy elasticity, which enables optimal budget allocation decisions [9, 17]. In this context, the various subsidy levels are the treatment, and the outcome is the consumer order placements. This outcome is quantified through the conversion rate, which is calculated as the ratio of the number of orders to the number of queries. Causal inference aims to estimate the treatment effect, also known as the uplift effect, in comparison to the control group. Ideally, the model should be trained on data from Randomized Control Trials (RCTs)[12, 13], where subsidies are randomly distributed to queries. However, RCTs are often impractical or too costly to conduct. Conversely, there is an abundance of observational data biased by the human-driven subsidy strategy, where confounding factors frequently exist in the input features, simultaneously influencing both the treatment and potential outcomes. For instance, subsidies are typically given to low-active or inactive users, which can mislead models because the conversion rate of consumers who receive subsidies is usually much lower than that of consumers who

do not receive subsidies in the observational data. Several methods have been devised to mitigate the bias introduced by confounders, where sensitivity analyses [15] fall short of providing precise point estimations, and instrumental variable approaches [1] only partially identify conditional average treatment effect. Moreover, most existing methods cannot handle multi-categorical or continuous subsidy levels.

With the evolution of online ride-hailing platforms, consumers are presented with a range of multi-class service options for the same trip. For example, Uber offers services like Comfort, UberX, and Uber Black for identical trips, each with varying costs and subsidy levels, analogous to multi-class airline seating. Given these diverse service and subsidy options, consumers may exhibit varying consumption preferences. When a consumer places an order for multiple services, the platform searches for service suppliers, and ultimately, only one service is provided to the consumer with a certain probability. In this paper, we introduce a multi-class ride-hailing subsidy system that has been deployed in a real-world environment. A **Multi-class Treatment Generalized Propensity Score Network** (MulTeNet) is introduced for modeling consumer subsidy elasticity, and a Model Predictive Control (MPC) framework is applied to optimize budget distribution.

2 The Proposed Method and Deployed System

Given a total financial budget, the objective is to identify an optimal policy for subsidizing queries to maximize overall revenue¹.

2.1 Multi-class Treatment Generalized Propensity Score Network

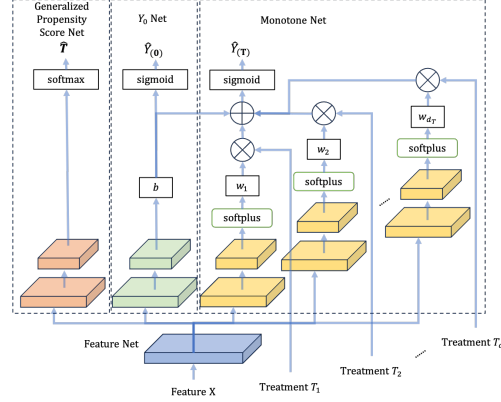


Fig. 1. The illustrated structure of the proposed MulTeNet

As shown in Fig. 1, the model is composed of three major components. Input features excluding treatments are processed by a common feature net, and the output is then directed to three distinct networks. The Generalized Propensity

¹ Due to limited space, the incentive for drivers is not discussed.

Score Net, which is a multi-class neural network, estimates the probability density function of treatment. The Y_0 Net takes the output from the feature net and estimates the base conversion rate without any treatment. The Monotone Net estimates the uplift effect of multi-class treatments. To ensure the monotonicity of the treatment effect, the softplus activation function is utilized. Inspired by [11], orthogonal regularization is introduced to rectify any biases present in the observational data.

2.2 Optimization Formulation

Similar to [17], the query-wise elasticity, as inferred by the proposed MulTeNet model, is then clustered by the origins, destinations and departure time. We formulate the budget distribution problem as follows:

$$\begin{aligned}
 & \max_{x_{i,j}} \sum_k \gamma_k \sum_i \hat{N}_i \hat{P}r_{k,i} \sum_j \hat{P}_{i,j} x_{i,j} \\
 \text{s.t.} \quad & \sum_k \gamma_k \sum_i \hat{N}_i \sum_j \hat{P}_{i,j} x_{i,j} c_{k,i,j} \leq B \\
 & 0 \leq \underline{u} \leq \sum_j \hat{P}_{i,j} x_{i,j} c_{k,i,j} \leq \bar{u}, \forall i, k \\
 & \sum_j x_{i,j} = 1, \forall i \\
 & x_{i,j} \in \{0, 1\}
 \end{aligned} \tag{1}$$

where k is the service class index, i is the cluster index, and j enumerates the multi-class treatment space. $x_{i,j}$ is a binary control variable indicating the j -th subsidy is allocated to cluster i . $\hat{P}_{i,j}$ is the estimated conversion rate for cluster i given the j -th treatment level, $\hat{P}r_{k,i}$ is the predicted order revenue for service class k in cluster i , $c_{k,i,j}$ is the financial cost of the j -th subsidy level, γ_k is the statistic probability that service class k is completed by suppliers, \hat{N}_i is the predicted order number, and B is the budget limit. The subsidy levels u are bounded by lower and upper limits \underline{u} and \bar{u} respectively. To efficiently solve the large-scale linear programming, a primal-dual approximation [10] is employed. Additionally, to account for modeling errors and randomness, MPC technology [14] is utilized at regular intervals with updated information.

2.3 System Implementation

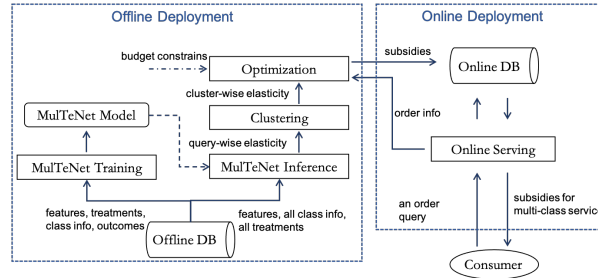


Fig. 2. The framework of our system

As depicted in Fig. 2, our system encompasses both offline and online components. The MulTeNet model is trained using historical features and treatment

information for different service classes. Subsequently, the model is deployed to infer query-wise elasticity at regular intervals. A clustering procedure is implemented to ensure consumer fairness. The cluster-wise elasticity and budget constraints are then fed into the optimization module, which generates a subsidy allocation dictionary for each service class. In the online phase, when a consumer inputs trip information, queries for different services are sent to the online server. The server simply looks up the dictionary and returns the corresponding results, significantly reducing the computational load in the online environment.

3 Empirical Experiments

Table 1. Casual inference results

Model	AUC	AUUC	QINI Coefficient
Yu&Etal2023	0.8065 ± 0.0001	1.8086± 0.0879	0.2155± 0.0023
DragonNet	0.8039± 0.0002	1.3825± 1.0462	0.2016± 0.0594
MulTeNet	0.7875± 0.0005	1.9363 ± 0.2135	0.3249 ± 0.0250

Empirical experiments were conducted to evaluate the performance of the proposed system. The dataset, comprising over 1.8 million queries with 200,000 subsidies, was extracted from the largest ride-hailing platform in mainland China between September and November 2023. For comparison, the baseline model Yu&Etal2023 employs a transfer learning approach [17]. Additionally, the DragonNet model, which is a multi-task DNN predicting potential outcomes and determining subsidy allocation, was also assessed. As shown in Table 1, our MulTeNet model achieves the highest AUUC (Area Under the Uplift Curve) and QINI coefficient, indicating its superior effectiveness in modeling uplift effects within the data. However, the AUC for MulTeNet is slightly lower, attributable to the introduction of treatment loss and regularization, which may impact response performance.

Table 2. Optimization results

	Yu&Etal2023	DragonNet	MulTeNet	No Subsidy
Revenue	1.0394	1.0283	1.0547	1.0
Orders	1.0224	1.0155	1.0298	1.0
Subsidy Rate	4.84%	4.88%	4.93%	-
ROI	0.79	0.56	1.05	-

The effectiveness of various subsidy Strategies was evaluated over a week in December 2023, with the target subsidy rate set to 5%. The 14-day historical data were used to forecast the subsequent 7 days, and the historical subsidy outcomes were fed back into the optimization solver daily to adjust the subsidy allocation for the remainder of the week. The ROI (Return On Investment) is defined as the increase in revenue over the total subsidy cost. As illustrated in Table 2, our system successfully controlled the subsidy rate to 4.93%, which is closer to the target level. More importantly, The total orders and revenue of the platform increased by 5.47% and 2.98%, respectively, and the ROI was 1.05, which is 34% higher than the best baseline.

References

1. Angrist, Joshua D and Imbens, Guido W and Rubin, Donald B: Identification of causal effects using instrumental variables. *Journal of the American statistical Association* **91**(434), 444–455 (1996)
2. Bimpikis, Kostas and Candogan, Ozan and Saban, Daniela: Spatial Pricing in Ride-Sharing Networks. *Operations Research* **67**(3), 744–769 (2019)
3. Castillo, Juan Camilo and Knoepfle, Dan and Weyl, Glen: Surge pricing solves the wild goose chase. In: *Proceedings of the 2017 ACM Conference on Economics and Computation*, pp. 241–242. (2017)
4. Chen, Le and Mislove, Alan and Wilson, Christo: Peeking beneath the hood of uber. In: *Proceedings of the 2015 internet measurement conference*, pp. 495–508. (2015)
5. Shu, Siqi and Chen, Zhengqi and Yu, Zhe and Cao, Shaosheng and Wu, Guobin and Shi, Donghai and Wang, Gaoang and Liu, Zuozhu and Chen, Xiqun and Na, Xiaoxiang and others: Modeling Freight-Sharing Platform Operations for Optimal Compensation Strategy Using Markov Decision Processes. In: *2022 IEEE 25th International Conference on Intelligent Transportation Systems (ITSC)*, pp. 1006–1011. (2022)
6. Wu, Tian and Zhang, Mengbo and Tian, Xin and Wang, Shouyang and Hua, Guowei: Spatial differentiation and network externality in pricing mechanism of online car hailing platform. *International Journal of Production Economics* **219**, 275–283 (2020)
7. Yan, Chiwei and Zhu, Helin and Korolko, Nikita and Woodard, Dawn: Dynamic pricing and matching in ride-hailing platforms. *Naval Research Logistics (NRL)* **67**(8), 705–724 (2020)
8. Yang, Jun and Wang, Wei and Dong, Yanshen and He, Xin and Jia, Li and Chen, Huan and Mao, Maoyu: GRFlift: uplift modeling for multi-treatment within GMV constraints **67**(8), 705–724 (2020)
9. Cohen, Peter and Hahn, Robert and Hall, Jonathan and Levitt, Steven and Metcalfe, Robert: Using big data to estimate consumer surplus: The case of uber, <http://www.nber.org/papers/w22627>. (2016)
10. Hao, Xiaotian and Peng, Zhaoqing and Ma, Yi and Wang, Guan and Jin, Junqi and Hao, Jianye and Chen, Shan and Bai, Rongquan and Xie, Mingzhou and Xu, Miao and others: Dynamic knapsack optimization towards efficient multi-channel sequential advertising. In: *International Conference on Machine Learning*, pp. 4060–4070. (2020)
11. Hatt, Tobias and Feuerriegel, Stefan: Estimating average treatment effects via orthogonal regularization. In: *Proceedings of the 30th ACM International Conference on Information & Knowledge Management*, pp. 680–689. (2021)
12. Kohavi, Ron and Deng, Alex and Frasca, Brian and Walker, Toby and Xu, Ya and Pohlmann, Nils: Online controlled experiments at large scale. In: *Proceedings of the 19th ACM SIGKDD international conference on Knowledge discovery and data mining*, pp. 1168–1176. (2013)
13. Kohavi, Ron and Longbotham, Roger and Sommerfield, Dan and Henne, Randal M: Controlled experiments on the web: survey and practical guide. *Data mining and knowledge discover* **18**, 140–181 (2009)
14. Qin, S Joe and Badgwell, Thomas A: A survey of industrial model predictive control technology. *Control engineering practice* **11**(7), 733–764 (2003)
15. Rosenbaum, Paul R and Rubin, Donald B: The central role of the propensity score in observational studies for causal effects. *Biometrika* **70**(1), 41–55 (1983)

16. Shi, Claudia and Blei, David and Veitch, Victor: Adapting neural networks for the estimation of treatment effects. *Advances in neural information processing systems* **32**, (2019)
17. Yu, Zhe and Xia, Chi and Cao, Shaosheng and Zhou, Lin and Huang, Haibin: A Consumer Compensation System in Ride-hailing Service. In: *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pp. 3240–3244. (2023)