BIRS workshop: Combinatorial Optimization for Online Platforms

Omar El Housni (Cornell University),
Srikanth Jagabathula (NYU Stern School of Business),
Mika Sumida (University of Southern California, Marshall School of Business),
Rajan Udwani (UC Berkeley)

April 5th–April 7th, 2024

1 Introduction

The workshop on combinatorial optimization was an insightful event that brought together experts from various fields to explore the latest advances and innovative paradigms in optimization. Combinatorial optimization, a branch of mathematics focused on finding the best solution from a discrete but extensive range of possibilities, is critical for online platforms such as e-commerce, social networks, ride-sharing, and streaming services. These platforms regularly confront complex combinatorial optimization challenges, including resource allocation, user matching, product recommendation, and item ranking. Discussions during the workshop centered on navigating large-scale problems and various uncertainties like customer demand, user behavior, substitution effects, and competitor actions. Emphasis was placed on balancing multiple objectives, such as revenue, profit, customer satisfaction, social welfare, and fairness. Participants explored scalable and effective algorithms for these complex and dynamic scenarios, highlighting the interdisciplinary nature of this research area, which involves Operations Research, Computer Science, Statistics, Economics, and Social Sciences. The workshop featured insightful talks by experts, interactive sessions where participants presented their work, and discussions on future directions. In the following sections, we provide an overview of the topics, discussions, and outcomes of this workshop.

2 Recent Developments and Key Findings

One of the primary applications of combinatorial optimization discussed extensively at the workshop is assortment optimization, which determines the best mix of products to offer to customers to maximize the firm's objective. Assortment optimization is a key concept for online platforms, as their processes often include customers making choices that directly impact their revenue. This field has been a key research avenue in revenue management, with various developments under multiple choice models such as the multinomial logit model (MNL) and the Markov Chain choice model. The talks presented significant contributions in the areas of revenue management, pricing, learning in assortment optimization, and discrete choice modeling.

Joint Learning and Assortment Optimization. Demand estimation is a crucial input to the assortment optimization problem, and these problems were often solved separately. This could lead to suboptimal decision-making, as better estimation of the demand model does not necessarily lead to better decisions. Prof. Antoine Desir presented a method to jointly estimate the demand model and make decisions, building on the literature of smart predict and optimize. Earlier literature used deep learning for demand estimation; however, Prof. Desir proposed adding a combinatorial layer to the demand model to learn optimal decisions more effectively.

Ningyuan Chen presented a multi-stage purchase model with learning under the MNL model. At each time step, the customer is offered an assortment, and the decision-maker can use the partial order induced by the customer choice in previous steps to make decisions in subsequent steps. They first focused on solving

the second-stage problem in the special case with two time periods, showing that this problem is NP-hard. Subsequently, they focused on the so-called *fixed policies*, which solve an assortment optimization problem at each time step independently of the previous time steps. They showed that fixed policies can recover a constant factor of the optimal solution. While this talk mainly highlighted the intractability of solving the considered model, it is one of the rare works in the assortment optimization literature to incorporate some form of demand learning in the decision-making process, opening avenues to further models with learning.

Algorithms for Assortment Optimization. The workshop highlighted several talks on static and dynamic assortment optimization problems and their interconnected nature with other problems arising in online platforms. For example, Xin Chen presented a model for multi-purchase assortment optimization with complementary effects, where platforms need to consider the complementary effects between several categories of products to optimize the selection. This leads to a rich class of challenging combinatorial problems that necessitate innovative frameworks. Vineet Goyal discussed threshold-based algorithms for online assortment platforms faced by e-retail, where they need to decide sequentially and irrevocably the selection of products to carry on the platform. This problem has strong connections to the prophet inequalities literature, leading to many discussions on the generalization of the setting presented by Prof. Goyal.

Ashwin Venkataraman presented their work on a joint pricing and assortment problem under a novel class of models, referred to as the mixture of boundary logits, which follows the consider-then-choose paradigm. In the boundary logit model, customers start by constructing a utility-maximizing consideration set, then make a second-stage choice from the consideration set according to the MNL model. The mixture of boundary logits considers multiple segments, each characterized by its own first and second-stage feature sensitivity. The authors showed that in the most general setting, the revenue maximization problem is hard, and they characterized the optimal solution in several special cases. One possible avenue of improvement is to randomize the first-stage choice, aligning it more with the standards of the assortment optimization literature.

Fairness in Assortment Optimization. The workshop also focused on incorporating fairness into assortment optimization, particularly in environments where decisions have significant socio-economic impacts. Traditionally, platforms aim to maximize revenue by selecting assortments that generate the highest returns, potentially leading to unfair exposure for some sellers. Prof. Ozge Sahin discussed how to incorporate fairness constraints into traditional assortment optimization problems to ensure that all sellers receive minimum market exposure based on factors like reputation and product quality. The study finds that the optimal solution involves randomizing up to n nested assortments, where n is the number of sellers, and can be determined in polynomial time. The implementation of fairness constraints enhances consumer welfare and benefits all parties, creating a win-win-win situation.

Fairness in decisions has become a crucial factor at various levels, and online platforms are often required or incentivized to take fair decisions towards their users. Two talks highlighted the notion of fairness in different areas. In vanilla assortment optimization models, only customers typically make choices on the platform, which can be unfair in contexts involving both a customer and a supplier, such as job matching platforms (UpWork, Fiverr) or dating apps (Bumble, Tinder). Two-sided assortment optimization models a platform consisting of both customers and suppliers, whose choices can both be captured by discrete choice models. Alfredo Torrico's talk considered multiple platform designs and analyzed the trade-offs between said designs, in addition to various approximation algorithms. An interesting avenue for future research is to consider more elaborate platform designs and more general choice models.

In addition to assortment optimization, the talks tackled other combinatorial optimization problems with fairness considerations. David Shmoys presented a joint work with Microsoft Research India, considering a fair version of the joint replenishment problem (JRP). This problem aims to coordinate the fulfillment of different orders utilizing the same resources to share the fixed cost associated with replenishment. The authors formulated this problem as an integer problem by extending the classic JRP and provided constant approximation algorithms.

Resource Allocation and Matching. Another innovative application discussed was the design of service menus for multi-class, multi-server queuing systems, where customers with diverse preferences choose among a variety of servers. The service provider's goal is to create a menu of service classes that achieves two conflicting objectives: maximizing the average matching reward of customers and minimizing the average waiting time of customers. Prof. Rene Caldentey explored this issue under conventional heavy traffic

conditions, where traffic intensity approaches unity from below. They showed that if the service provider concentrates only on reducing average delay or enhancing total matching reward, very simple menus prove to be optimal. Moreover, they introduced Mixed Integer Linear Programming (MILP) formulations to fine-tune the delay-reward trade-off across a spectrum of functional and pertinent menu families.

Ali Aouad presented recent work on advancements in dynamic matching, where they introduced methods for designing tighter LP relaxations with adaptive matching policies that can be approximated in polynomial time. Varun Gupta discussed a finite horizon online resource allocation/matching problem that aims to combine resources from a finite set of types into feasible configurations. They showed that a simple greedy algorithm is efficient and can bound regret in some scenarios. Rad Niazadeh considered a dynamic staffing problem, as part of a joint project with Amazon, proposing a minimax optimal online algorithm for their problem.

Advanced models in network revenue management (NRM) address a complex combinatorial problem where products, composed of resources, are sold to stochastically arriving customers. Prof. Will Ma discussed a randomized rounding approach to NRM, motivated by developments in Online Contention Resolution Schemes (OCRS), contrasting with traditional methods that use deterministic linear programming or approximate dynamic programming. These models are crucial for handling scenarios with customer substitutions and varying demand correlations, typical in dynamic online auction environments. They provided a unifying framework for applying OCRS to these challenges, delineating the impact of substitution, and establishing a distinction between the guarantees achievable with and without substitution under general resource constraints.

Finally, Prof. Jackie Baek explored a real-world application in collaboration with Keheala, an organization in Kenya that offers a digital service to improve medication adherence among tuberculosis (TB) patients. The organization tackles the challenge of optimizing personalized interventions to maximize long-term results in a context where interventions are costly and capacity-constrained. They introduced a novel approach that approximates one step of policy iteration and simplifies the process to a straightforward prediction task using available data, eliminating the need for online experimentation. They further derived theoretical guarantees for this approach based on techniques from combinatorial optimization. This algorithm is generic, model-free, and operates effectively regardless of the underlying patient behavior model. They also presented evidence of its robust performance, both theoretically and empirically.

Panel discussion. The panel discussion titled "Theory and Practice of Algorithmic Revenue Management," moderated by Srikanth Jagabathula and featuring panelists David Shmoys, Rene Caldentey, and Xin Chen delved into issues surrounding the evolution of revenue management as a field. Key points included concerns about revenue management's growing divergence from practice, how best to position and evaluate theoretical research in revenue management, potential under-appreciation of empirical research within the field, and the increasing demand for computational resources, including comprehensive datasets, to keep RM aligned with practice.

The panel highlighted the importance of creating a standardized corpus of datasets for training and benchmarking of revenue management algorithms. Drawing a parallel to the transformative impact of the ImageNet dataset in the computer science community, the panel discussed how a similar initiative could spur research as well as attract collaborations with researchers in adjacent fields such as computer science.

3 Future Directions and Conclusion

The talks and discussions highlighted several challenges and future research directions, such as exploring the potential to utilize deep learning methods with combinatorial layers to bring new insights even to classic well-studied problems, the exploration of service system design in the context of congestion and strategic customer behavior, developing models for capturing and optimizing the effect of complementary product categories on revenue, exploring the role of adaptivity in two-sided online markets and new models and methodologies in assortment optimization and online resource allocation that hold significant promise for the advancement of the field.

In conclusion, the workshop not only underscored the critical role of combinatorial optimization in transforming online platforms but also set a robust foundation for future research directions. It also brought different parts of the community together and garnered interactions on multi-disciplinary topics. We believe

that the interactions initiated between participants during this workshop have the potential to lead to new groundbreaking work on these challenging problems.

References

[1] A.M. Farzaneh, S. Modaresi, and A. Venkataraman, *Available at SSRN* **4357465** (2023).

[2] N. Chen, P. Gao, C. Wang and Y. Wang, Assortment optimization for the multinomial logit model with repeated customer interactions, *Available at SSRN* **4526247** (2023).

[3] O. El Housni, A. Torrico and U. Hennebelle, Two-sided Assortment Optimization: Adaptivity Gaps and Approximation Algorithms, *arXiv preprint arXiv:2403.08929* **4526247** (2024).

[4] V. Suriyanarayana, V. Sivashankar, S. Gollapudi and David B. Shmoys, Improved Approximation Algorithms for the Joint Replenishment Problem with Outliers, and with Fairness Constraints. In *Proceedings of the 2024 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)* **2793-2828** (2024).

[5] Y. Feng, V. Manshadi, R. Niazadeh, and S. Neyshabouri, Robust Dynamic Staffing with Predictions. *Available at SSRN* (2023).

[6] K. Talluri, G. Van Ryzin, Revenue management under a general discrete choice model of consumer behavior. *Management Science* **50(1)** (2004).

[7] P. Rusmevichientong, H. Topaloglu, Robust assortment optimization in revenue management under the multinomial logit choice model. *Operations research* **60(4)** (2012).

[8] A. Désir, V. Goyal, D. Segev, C. Ye, Constrained assortment optimization under the markov chain–based choice model. *Management Science* **66(2)** (2020).

[9] A. Désir, V. Goyal, B. Jiang, T. Xie, J. Zhang, Robust assortment optimization under the markov chain choice model. *Operations Research* (2023).

[10] A. Aouad, V. Farias, R. Levi, Assortment optimization under consider-then-choose choice models. *Management Science* **67(6)**(2021).

[11] V. Goyal, S. Humair, O. Papadigenopoulos, and A. Zeevi, MNL-Prophet: Sequential Assortment Selection Under Uncertainty.' *arXiv preprint arXiv:2308.05207*, 2023.

[12] V. Gupta, Greedy Algorithm for Multiway Matching with Bounded Regret. *Operations Research*, INFORMS, 2022.

[13] X. Chen, J. Li, M. Li, T. Zhao, and Y. Zhou, Assortment Optimization under the Multivariate MNL Model. *arXiv preprint arXiv:2209.15220*, 2022.

[14] A. N. Elmachtoub and P. Grigas, Smart "predict, then optimize", *Management Science* **68**(1), 9-26 (2022).

[15] A. Aouad and A. Désir, Representing random utility choice models with neural networks, *arXiv preprint* arXiv:2207.12877 (2022).

[16] L. Wentao, O. Sahin, and R. Wang, A simple way towards fair assortment planning: Algorithms and welfare implications, *Available at SSRN* 4514495 (2023).

[17] R. Caldentey, L. A. Hillas, and V. Gupta, Designing service menus for bipartite queueing systems, *Operations Research* (2024).

[18] W. Ma, C. MacRury, and J. Zhang, Online contention resolution schemes for network revenue management and combinatorial auctions, *arXiv preprint* arXiv:2403.05378 (2024).

[19] J. Baek, et al., Policy optimization for personalized interventions in behavioral health, *arXiv preprint* arXiv:2303.12206 (2023).