Democratizing Travel Personalization via Central Recommendation Platform

Chana Ross[1], Tomer Ovadia[1], Jake Mooney[1], Amit Meitin[1], Eytan Kabilou[1], Mush Kabalo[1] and Dmitri Goldenberg[1]

[1]*Booking.com, Tel Aviv, Israel*

Abstract

Recommender systems play a crucial role in e-commerce platforms, reducing the information overload problem by providing and prioritizing relevant information based on a user's implicit and explicit preferences. Large e-commerce platforms often rely on a multitude of machine learning models at the same time to optimize the many aspects of the site that can be improved. Occasionally there are duplicate recommender systems that solve a similar or even the same recommendation task, a situation often caused by evolving business objectives, use-case constraints or even just a lack of synchronicity between teams of a large organization.

In this work we demonstrate a centralized Recommendation Platform at Booking.com - one of the world's leading online travel platforms. The system is created to reduce the duplication of work, provide utilities for authors of new recommendation models, increase impact by accelerating adoption of better solutions to common recommendation problems, and serve as a single, trusted point of access for recommendations across the website. It allows us to democratize the usage of recommendations across the company and ease the development of new sophisticated models. This, in turn, allows standardization of recommendations tasks and increases the adoption of recommendation systems by various product teams, to bring a personalized experience to each of our customers.

Keywords

Personalization, Travel, Recommender Systems

1. Introduction

Online travel platforms offer a vast variety of products that serve different needs of travelers. Such platforms often rely on recommender systems to assist various customer decisions, seeking to help narrow down diverse offerings by suggesting well-fit options in a personalized manner [1]. Such platforms often rely on a cascade of different recommender systems [2] resolving various personalization tasks across the customer journey [3].

Planning a trip is a particularly complex task requiring important and often expensive decisions by the customers, within limited time and information. Availability, information gaps, budget, timing, preferences and even weather constraints introduce non-negligible complexity

RecSys Workshop on Recommenders in Tourism (RecTour 2022), September 22th, 2022, co-located with the 16th ACM Conference on Recommender Systems, Seattle, WA, USA

chana.ross@booking.com (C. Ross); tomer.ovadia@booking.com (T. Ovadia); jacob.mooney@booking.com (J. Mooney); amit.meitin@booking.com (A. Meitin); Eytan.kabilou@booking.com (E. Kabilou); mush.kabalo@booking.com (M. Kabalo); dima.goldenberg@booking.com (D. Goldenberg)

[4]. While some travelers have an exact plan in mind, others have some degree of open-mindedness. Customers are likely to benefit from travel recommendations differently in different contexts. Because customers are not used to expressing their travel intent explicitly when interacting with online travel platforms, recommender systems need to be able to interpret a full spectrum using implicit cues as inputs, including cues set by user context. Therefore, the work on recommendation tasks often involves multiple teams in the same company, developing complementary and even parallel models to resolves various recommendations tasks. This causes a need in a centralized platform for recommendations to allow fast development of recommender systems on the one hand, and democratize their usage across multiple product teams on the other hand.

In this work we present the *Recommendation Platform* at Booking.com, serving various recommendation tasks within a single system. We demonstrate the system through the use-case of destinations recommendations [5] which requires a complex consideration of available context data [6], customer needs and recommendation purpose. Recommending travel destinations presents unique challenges including, but not limited to, continuous cold-start [7], seasonality, timing of recommendation and delayed feedback [8]. These challenges necessitate employing sophisticated modeling approaches, such as sequential learning [9], combining different feature types within the model [10], and introducing online adaptive algorithms. We present the system design and the work flow of our centralized Recommendation Platform and explain its benefits in the multi-contextual use cases of destination recommendations at Booking.com.

2. System Overview

2.1. Centralized Platform

Similarly to other data-driven centralized platforms at Booking.com, such as the experimentation [11] and the machine learning [2] platforms, our goal is to simplify both the development and the consuming process of new and existing recommender systems in the company. The platform allows to quickly integrate new models and provide the needed input features, while at the same time provides a simple API to customer facing systems, that consume these recommendations.

The Recommendation platform is a central ML-driven system for enabling a personalized customer experience from one source, recommending destinations, properties, attractions and other trip products to another. The system is able to serve the customer on all possible connection points such as web, mobile or direct marketing. This ensures consistent and relevant recommendations across products and customer journey moments. Those recommendations are expected to be relevant, engaging and insightful based on the customer segment preferences, search and reservations history and the customer stage in the booking funnel. In addition, the platform provides a set of tools and services for reducing the friction of recommendation models development and enriching models with relevant features. Models are monitored, continuously optimized, and compared with one another for maximizing business impact. Recommendations and enrichment features are centrally accessible to consume at scale.

2.2. Architecture

A standard flow of the system (Figure 1) will start with the customer accessing a relevant page (whether it is on a web platform or a mobile device) designed to display recommendations. At this point an HTTP request will be sent to the *Recommendation Platform* via our GraphQL layer or alternatively via REST API which acts as the main Gateway. The request will include context about the intent and recommendation type, user, platform, tracking, geo data etc. From there a central module in the system will manage the significant decision making for what type of recommendations should be retrieved and then for gathering the information and making it accessible back to the customer. First the system fills in some additional information about the user (User Profile Enrichment, complying with applicable privacy laws and regulations when collecting and using this information). Once the user information is available, the system chooses a model which "best" fits the current use-case. The term "best" here is relative to each use-case and is dependent on required performance, latency and what available information exists at inference time. In most cases, the further down the funnel a customer is, the more information we have at our disposal to feed into the model chosen. Our platform accesses the relevant model with the required payload and returns the top N destinations (these are chosen based on a matching score the model gives each item and customer).

The system is integrated with the relevant supply, ML serving [2], experimentation [11], pricing, and content services. After getting the list we enrich the needed content about the destination (Destination Info Enrichment), such as: pricing/visual image/translated content/etc, and eliminate options we find as less relevant according to some business rules. Once we have the final list and the content associated with it the system sends this content back in an API response to the user in order for them to take care of rendering the results for the client.

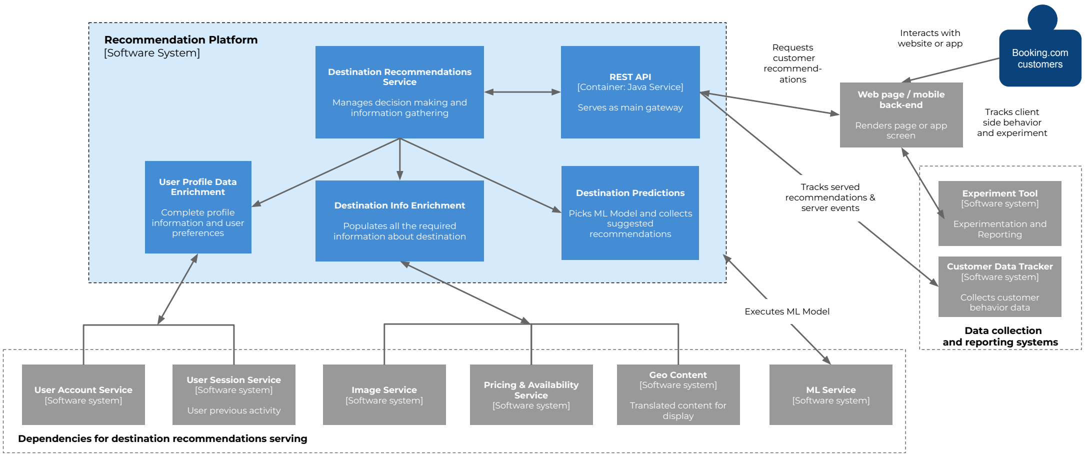

Figure 1: Schematic Diagram of Recommendation Platform flow and components

3. Destination recommendations Use-Case

Destination recommendations are one of the key use-cases for the recommendation platform [5]. We consider a set of recommender systems distributed across our website, each of which is tailored to specific stages of the traveller's booking journey. They can be described by a funnel-like structure as follows:

- **Inspirational**: The goal of these recommendations is to trigger *inspirational* recommendations to travelers who have expressed no intent at all, such as via email campaigns to customers who signed up in the past.
- **Cold Start**: Travelers visiting our platform without showing intent on where to travel. These recommendations are displayed in the index page of our web pages and mobile apps.
- **Alternative**: For travelers browsing a specific destination, our platform can offer *alternative* options. There are several flavours such as *Nearby Destinations* aiming at expanding the availability of the current choice and *Similar Destinations* aiming to offer an alternative plan with a similar experience.
- **Complementary**: After completing a reservation in a specific destination, our system recommends *the next* destination to visit in order to extend the trip.

The recommendations are powered by more than 50 machine learning models, sometimes simultaneously. The algorithms used to power the use case, could differ according to the task: The variety of solutions ranges from simple lookup-table models context-aware recommenders, sequential models [9, 10, 12] and online exploration algorithms.

4. Playground

The platform provides an interactive playground (see Figure 2) for teams within Booking.com to experience and get a "feel" for the different models the system has to offer, based on their chosen inputs. This playground can be used by all types of recommendation consumers (machine learning scientists, developers and even product managers trying to understand the algorithm). In a case of destination recommendations, a practitioner can quickly test various models and feature inputs, to perform a "sanity check" and test the model fitness for the selected use-case.

5. Model Evaluator

While the playground functionality allows a quick sanity check of models, and make the model selection process accessible for non-technical practitioners, it may bias the selection process towards cherry-picked and un-generalized results. Therefore, as a part of the model comparison suite, the recommendation platform provides a model evaluation tool. This component allows to compare the empirical performance (such as Accuracy @ TopK) of several models, given the same set of features and selected input traffic. Such comparison allows to quickly generate candidates for online experimentation, while at the same time performs a seasonal monitoring of existing deployed solutions.

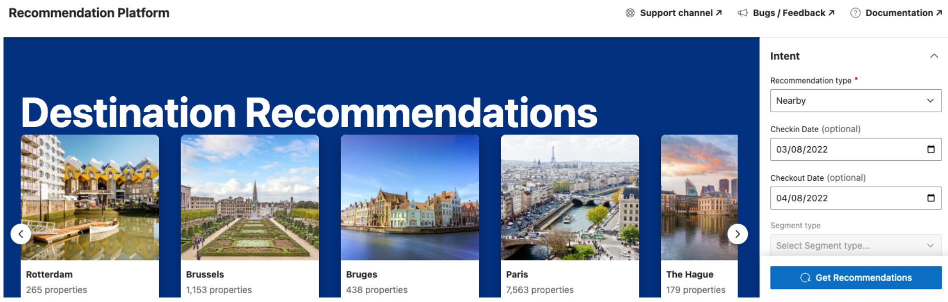

Figure 2: A screenshot of the Recommendation Platform playground

Figure 3 demonstrates an example comparison of three destination recommendations models in an "Alternative destination" scenario. The charts plots the Accuracy @ Top 10 performance of a personalized model, compared to simple "Nearby" and "Domestic" recommends. It is worth mentioning that all three models are already deployed in several pages across the website. We observe that the personalized model is losing its predictability during COVID-19 period and recovering back to the leading position afterwards. At the same time, the two "basic models" demonstrate a relatively consistent performance. Moreover, the two models are providing practically the same recommendations, suggesting a valid consideration to decommission one of them and reduce redundancy.

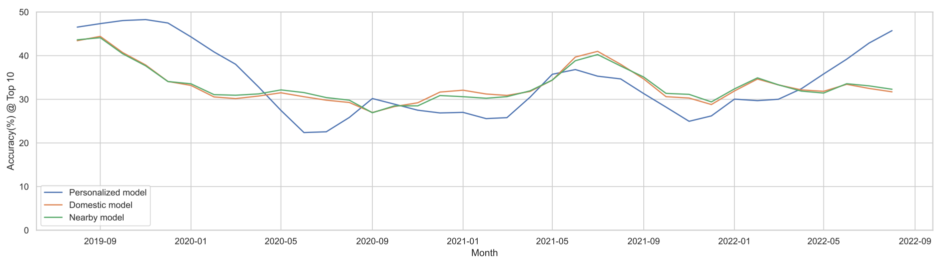

Figure 3: Demonstration of three destination recommendations models performance during COVID-19

6. Discussion

Our paper demonstrates the main architecture used in Booking.com for destinations recommendations. In addition, we show how this system is displayed in a playground allowing others in the company to investigate their models and better understand the results and enhanced information received from the platform. The system can be used during and after the model is built for testing and building future models. The main benefit of our platform is the comparison and flexibility it provides scientists in the company to test their model throughout the website and user journey. Unlike most models which need specific work for each entry point in the website, we enhance the features needed for the model based on the context of the user entered (number of rooms for example) and the searched destination.

References

[1] D. Goldenberg, K. Kofman, J. Albert, S. Mizrachi, A. Horowitz, I. Teinemaa, Personalization in practice: Methods and applications, in: Proceedings of the 14th ACM International Conference on Web Search and Data Mining, WSDM '21, Association for Computing Machinery, New York, NY, USA, 2021, p. 1123–1126. URL: https://doi.org/10.1145/3437963.3441657. doi:10.1145/3437963.3441657.

[2] L. Bernardi, T. Mavridis, P. Estevez, 150 successful machine learning models: 6 lessons learned at booking.com, in: Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, 2019, pp. 1743–1751.

[3] I. Kangas, M. Schwoerer, L. J. Bernardi, Recommender systems for personalized user experience: Lessons learned at booking. com, in: Fifteenth ACM Conference on Recommender Systems, 2021, pp. 583–586.

[4] T. H. N. Do, W. Shih, Destination decision-making process based on a hybrid mcdm model combining dematel and anp: The case of vietnam as a destination, Modern Economy 7 (2016) 966–983.

[5] D. Goldenberg, S. Mizrachi, A. Horowitz, I. Kangas, M. Schwoerer, A. Mozzato, M. Ferretti, K. Panagiotis, L. Bernardi, I know what you did next summer: Challenges in travel destination recommendation (2021).

[6] J. Kiseleva, M. J. Mueller, L. Bernardi, C. Davis, I. Kovacek, M. Stafseng Einarsen, J. Kamps, A. Tuzhilin, D. Hiemstra, Where to go on your next trip? optimizing travel destinations based on user preferences, in: Proceedings of the 38th International ACM SIGIR Conference on Research and Development in Information Retrieval, 2015, pp. 1097–1100.

[7] L. Bernardi, J. Kamps, J. Kiseleva, M. J. Müller, The continuous cold start problem in e-commerce recommender systems, 2015.

[8] T. Mavridis, S. Hausl, A. Mende, R. Pagano, Beyond algorithms: Ranking at scale at booking. com, in: Proceedings of the Fourth Workshop on Recommendation in Complex Scenarios. CEUR-WS, 2020.

[9] D. Goldenberg, P. Levin, Booking.com multi-destination trips dataset, in: Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '21, Association for Computing Machinery, New York, NY, USA, 2021, p. 2457–2462. URL: https://doi.org/10.1145/3404835.3463240. doi:10.1145/3404835.3463240.

[10] S. Mizrachi, P. Levin, Combining context features in sequence-aware recommender systems., in: Late-Breaking Results of the 13th ACM Conference on Recommender Systems, RecSys '19, 2019.

[11] R. L. Kaufman, J. Pitchforth, L. Vermeer, Democratizing online controlled experiments at booking. com, arXiv preprint arXiv:1710.08217 (2017).

[12] B. Schifferer, C. Deotte, J.-F. Puget, G. d. S. Pereira Moreira, G. Titericz, J. Liu, R. Ak, Using deep learning to win the booking.com wsdm webtour21 challenge on sequential recommendations, in: Proceedings of the ACM WSDM Workshop on Web Tourism (WSDM Webtour '21), 2021.