RecSys Challenge 2019: Session-based Hotel Recommendations

Peter Knees
TU Wien
Vienna, Austria
peter.knees@tuwien.ac.at

Yashar Deldjoo
Polytechnic University of Bari
Bari, Italy
yashar.deldjoo@poliba.it

Farshad Bakhshandegan Moghaddam
Karlsruhe Institute of Technology
Karlsruhe, Germany
uwowi@student.kit.edu

Jens Adamczak
trivago N.V.
Düsseldorf, Germany
jens.adamczak@trivago.com

Gerard-Paul Leyson
trivago N.V.
Düsseldorf, Germany
gerard-paul.leyson@trivago.com

Philipp Monreal
trivago N.V.
Düsseldorf, Germany
philipp.monreal@trivago.com

ABSTRACT

The workshop features presentations of accepted contributions to the RecSys Challenge 2019 organized by trivago, TU Wien, Politecnico di Bari, and Karlsruhe Institute of Technology. In the challenge, which originates from the domain of online travel recommender systems, participants had to build a click-prediction model based on user session interactions. Predictions were submitted in the form of a list of suggested accommodations and evaluated on an offline data set that contained the information what accommodation was clicked in the later part of a session. The data set contains anonymized information about almost 16 million session interactions of over 700.000 users visiting the trivago website.

The challenge was well received with 1509 teams that signed up and 607 teams teams that submitted a valid solution. 3452 solutions were submitted during the course of the challenge.

CCS CONCEPTS

• Information systems → Recommender systems; **• Applied computing →** *E-commerce infrastructure.*

KEYWORDS

recommender systems; challenge; travel meta-search; session-based; context-aware; accommodation recommendation; benchmark

ACM Reference Format:

Peter Knees, Yashar Deldjoo, Farshad Bakhshandegan Moghaddam, Jens Adamczak, Gerard-Paul Leyson, and Philipp Monreal. 2019. RecSys Challenge 2019: Session-based Hotel Recommendations. In *Thirteenth ACM Conference on Recommender Systems (RecSys '19), September 16–20, 2019, Copenhagen, Denmark.* ACM, New York, NY, USA, 2 pages. https://doi.org/10.1145/3298689.3346974

1 INTRODUCTION

The annual ACM Recommender System Challenge, every year, introduces a unique recommendation problem to academia and industry. As in the previous years, e.g. [1, 3], the goal of the challenge is to develop novel approaches and apply predictive techniques on data sets that originate from real world applications.

For the RecSys Challenge 2019,[1] real-world data was provided by industry partner trivago.[2] trivago is a global hotel search platform providing travelers with aggregated information about the characteristics of accommodations to help them make an informed decision and finding their ideal place to stay at the lowest rate. For the first time, the RecSys Challenge therefore dealt with a real-world task from the online travel domain, which is a particularly interesting and challenging field for recommender systems due to multiple stakeholders and value-awareness in recommendations, sparsity of user data, variety in accommodations, extreme cold-start scenarios, highly dynamic search criteria and computational requirements on how fast results need to be delivered for a good user experience [2].

2 DATA

The data set provided for the purpose of the challenge contains anonymized information about almost 16 million session interactions of over 700.000 users visiting the trivago website. This data consists of a training and test set, and metadata for accommodations (items).The training and test data sets consists of 12 columns with information about users and the session interactions.

- **user_id:** identifier of the user
- **session_id:** identifier of each session
- **timestamp:** UNIX timestamp for the time of the interaction
- **step:** step in the sequence of actions within the session
- **action_type:** identifier of the action taken by the user
- **reference:** reference value of the action as described for the different action types
- **platform:** country platform used for the search
- **city:** name of the current city of the search context
- **device:** device used for the search
- **current_filters:** list of filters active at given timestamp
- **impressions:** list of items that were displayed to the user
- **prices:** list of prices of the impressions

Additional metadata for each accommodation was specified in the form of a list of filter options that are applicable for each item.

RecSys '19, September 16–20, 2019, Copenhagen, Denmark

ACM ISBN 978-1-4503-6243-6/19/09.
https://doi.org/10.1145/3298689.3346974

[1] http://www.recsyschallenge.com/2019/
[2] https://www.trivago.com

3 CHALLENGE

The problem statement of the challenge required participants to build models based on user actions up to a specified time (split date) and had to predict the accommodations that have been clicked out in later sessions after the split date. Participants were provided with the type of action (e.g. filter usage, search refinements, item interactions) and the context of the click-out (list of impressed items with prices).

Participants had to provide a list of maximum 25 items for each click-out ordered by preferences for the specific user. Submitted lists were evaluated using the Mean Reciprocal Rank metric, calculated as $MRR = \frac{1}{N}\sum_{i=1}^{N}\frac{1}{\text{rank}_i}$, where rank_i is the rank position of the actually clicked-out item within the i^{th} list and N the number of evaluated lists. MRR thus uses the position of the clicked item as a proxy for relevance, i.e. the higher the actually clicked item appears on the list, the higher the score.

The specifics of the problem statement required participants to address a series of challenges. They had to deal with a data set of diverse and time-dependent information that required to identify meaningful information in the search history of the user, the context of the click-out, or both, and build models that are capable in incorporating both. Most importantly, participants had to overcome the extreme cold-start scenario in which little information was provided on an individual user level and had to find ways to effectively use the information on session or aggregated user level.

4 PARTICIPATION

Participants actively submitted solutions throughout the challenge in line with the submission guidelines that limited to 1 submission every 12 hours. On average a team submitted 5.3 solutions, the most active team submitted 119 solutions.

Figure 1 illustrates the distribution of submissions per team. The bulk of the teams submitted few submission but there are also very active participants submitting over 100 solutions.

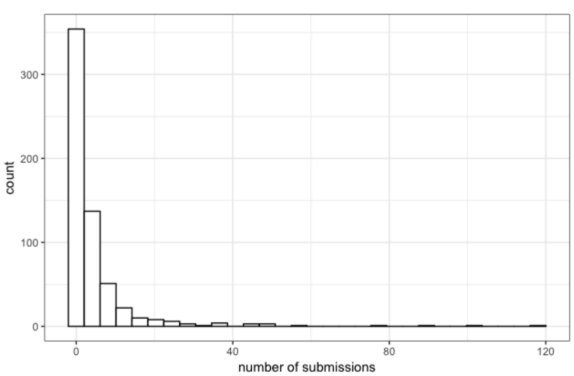

Figure 1: Number of submissions per team.

Figure 2 shows the number of submissions per day. Activity was high throughout the challenge and spiked in the last days of the challenge when participants refined their solutions.

To the accompanying RecSys Challenge 2019 Workshop, 13 papers were submitted, of which 9 were accepted for presentation after review by the workshop program committee.

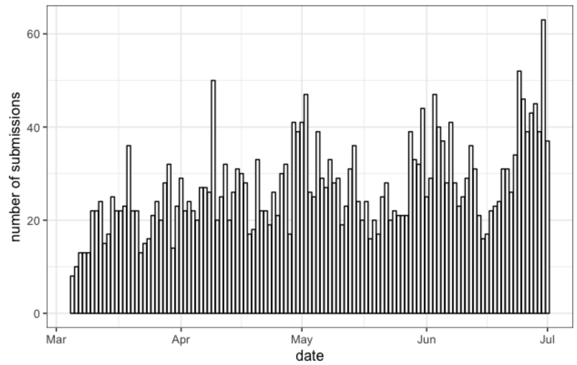

Figure 2: Number of submissions per day.

5 CHALLENGE RESULTS

The maximum MRR achieved on the public test set was 0.689. The team with the highest score on the public test set also achieved the highest score on the private test set (0.686) indicating that the solution was not prone to overfitting and the public test set a representative sample.

Figure 3 highlights the development of the maximum MRR score for each day of the challenge. At the beginning of the challenge the scores saw a steep increase as participants realized characteristics of the data set that were promising indicators for the click behavior. After the initial jump, the scores continuously increased until the end of the challenge.

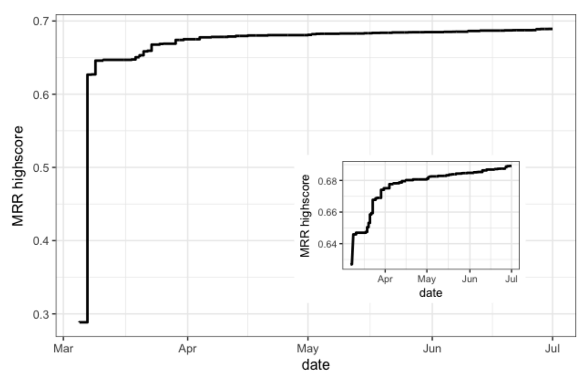

Figure 3: Development of the highest MRR score throughout the challenge.

REFERENCES

[1] Fabian Abel, Yashar Deldjoo, Mehdi Elahi, and Daniel Kohlsdorf. 2017. RecSys Challenge 2017: Offline and Online Evaluation. In *Proceedings of the Eleventh ACM Conference on Recommender Systems, RecSys 2017, Como, Italy, August 27-31, 2017.* 372–373. https://doi.org/10.1145/3109859.3109954

[2] Jens Adamczak, Gerard-Paul Leyson, Peter Knees, Yashar Deldjoo, Farshad Bakhshandegan Moghaddam, Julia Neidhardt, Wolfgang Wörndl, and Philipp Monreal. 2019. Session-Based Hotel Recommendations: Challenges and Future Directions. arXiv:1908.00071

[3] Ching-Wei Chen, Paul Lamere, Markus Schedl, and Hamed Zamani. 2018. Recsys challenge 2018: automatic music playlist continuation. In *Proceedings of the 12th ACM Conference on Recommender Systems, RecSys 2018, Vancouver, BC, Canada, October 2-7, 2018.* 527–528. https://doi.org/10.1145/3240323.3240342