Challenges of Session-Aware Recommendation in E-Commerce (Keynote)

Dietmar Jannach
TU Dortmund, Germany
dietmar.jannach@tu-dortmund.de

KEYWORDS

Recommender systems, Session-aware Recommendation

1 MOTIVATION

Research in the field of recommender systems is in many cases based on the matrix completion problem abstraction. While being able to assess the user's *general* preferences towards individual items is important, this popular problem abstraction often cannot fully capture certain aspects that are important for the success of a recommender in practice, in particular in e-commerce settings [3].

(1) First, in e-commerce scenarios users often visit an online shop with a very specific shopping intent, and a recommender system, to be successful, must be able to adapt its recommendations to the particular contextual situation and short-term preferences of the user.
(2) Second, in some cases it might be relevant to know which items the user inspected in his or her last session, and a recommender could use this knowledge to *remind* the user of such items.
(3) Third, in some domains, also aside from e-commerce, considering popularity trends in the user community could be helpful when deciding on what to recommend to users.
(4) Fourth, some users might be interested in certain items only in case they are currently discounted. For such users, recommending items that are on sale could be a promising strategy.

To be able to analyze research questions like these, a different problem abstraction is required. Instead of a user-item rating matrix, the input to the recommendation problem is rather a time-ordered sequence (log) of user actions of different types, e.g., an item view event, a purchase, etc. Correspondingly, other computational tasks than rating prediction have to be considered, including the prediction of the next user action, the identification of trends, or the consideration of sequence constraints.

2 CONTENTS OF THE TALK

2.1 Defining Sequence-Aware Recommendation

We will first highlight the importance of what we call "sequence-aware" recommender systems in practice and categorize different recommendation scenarios where the main input to the problem is a time-ordered series of logged user actions. We will consider different computational tasks in that context and then specifically focus on session-*based* recommendation problems (where only the interactions of the current session are known) and session-*aware* ones (where we also know previous sessions of the current user).

2.2 E-Commerce Case Studies

In the remainder of the talk, we will focus on the specific problems mentioned in the introduction and present insights from recent research works. The case studies are based on a large data set containing logged user interactions of a major European fashion retailer.

2.2.1 Considering short- and long-term interests. In the first case study [1], we compare the performance of different heuristic approaches to adapt the system's recommendations to the estimated short-term interests of the user. The results show that while the choice of the underlying long-term model is relevant, considering short-term interests in the right way has much more impact on recommendation accuracy.

2.2.2 On the value of reminders. Since the first study revealed that reminding users of things they have inspected (but not purchased) in the recent past can be an effective strategy, we then explore more elaborate reminding techniques than just presenting the recently viewed items in reverse chronological order [4].

2.2.3 Deriving recommendation success factors from log data. Since the available log data also contains information about which items were recommended to users and which of these items they actually inspected, we then re-construct in a systematic way which factors contributed to the success of the presented recommendations. The analysis shows that besides the consideration of the recent interests, recommendations are particularly successful when they are related to currently trending or to discounted items [2].

2.2.4 Operationalizing the insights into algorithms. Finally, we present a novel algorithmic approach to predict the next user interaction that considers all of the above-mentioned factors (short-term intents, reminders, trends and discounts) in an integrated way. We frame the recommendation problem as a classification task and our experiments show that a deep neural network leads to a better performance than when using Random Forests or when using a weighted hybrid scoring approach.

2.3 Outlook

The talk ends with a discussion of open questions and possible future directions in the field.

ComplexRec 2017, Como, Italy.

REFERENCES

[1] Dietmar Jannach, Lukas Lerche, and Michael Jugovac. 2015. Adaptation and Evaluation of Recommendations for Short-term Shopping Goals. In *Proc. RecSys '15*. 211–218.

[2] Dietmar Jannach and Malte Ludewig. 2017. Determining Characteristics of Successful Recommendations from Log Data – A Case Study. In *Proc. SAC '17*.

[3] Dietmar Jannach, Paul Resnick, Alexander Tuzhilin, and Markus Zanker. 2016. Recommender Systems - Beyond Matrix Completion. *Commun. ACM* 59, 11 (2016), 94–102. DOI: http://dx.doi.org/10.1145/2891406

[4] L. Lerche, D. Jannach, and M. Ludewig. 2016. On the Value of Reminders within E-Commerce Recommendations. In *Proc. UMAP '16*. 27–25.