Keynote: Session-based Recommendation – Challenges and Recent Advances

Dietmar Jannach

AAU Klagenfurt, 9020 Klagenfurt
dietmar.jannach@aau.at

Abstract. In many applications of recommender systems, the system's suggestions cannot be based on individual long-term preference profiles, because a large fraction of the user population are either first-time users or returning users who are not logged in when they use the service. Instead, the recommendations have to be determined based on the observed short-term behavior of the users during an ongoing session. Due to the high practical relevance of such session-based recommendation scenarios, different proposals were made in recent years to deal with the particular challenges of the problem setting.
In this talk, we will first characterize the session-based recommendation problem and its position within the family of sequence-aware recommendation. Then, we will review algorithmic proposals for next-item prediction in the context of an ongoing user session and report the results of a recent in-depth comparative evaluation. The evaluation, to some surprise, reveals that conceptually simple prediction schemes are often able to outperform more advanced techniques based on deep learning. In the final part of the talk, we will focus on the e-commerce domain. We will report recent insights regarding the consideration of short-term user intents, the importance of considering community trends, the role of reminders, and the recommendation of discounted items.

Keywords: Recommender Systems · Session-based Recommendation

1 Introduction

Recommender systems (RS) are tools that help users find items of interest within large collections of objects. They are omnipresent in today's online world, and many online sites nowadays feature functionalities like Amazon's "Customers who bought ... also bought" recommendations.

Historically, the recommendation problem is often abstracted to a matrix-completion task, see [8] for a brief historical overview. In such a setting, the goal is to make preference or rating predictions, given a set of preference statements of users toward items. These statements are usually collected over longer periods of time. In many real-world applications, however, such long-term profiles often do not exist or cannot be used because website visitors are first-time users, are not logged in, or take measures to avoid system-side tracking. These scenarios lead to what is often termed a *session-based recommendation* problem in the literature. The specific problem in these scenarios therefore is to make helpful recommendations based only on information derived from the ongoing session, i.e., from a very limited set of recent user interactions.

While the matrix completion problem formulation still dominates the academic research landscape, in recent years, increasing research interest can be observed for session-based recommendation problems. This interest is increased not only due to the high practical relevance of the problem, but also due to the availability of new research datasets and the recent development of sophisticated prediction models based on deep neural networks [2, 3, 13].

In this talk, we will first characterize session-based recommendation problems as part of the more general family of *sequence-aware* recommendation tasks. Next, we will briefly review existing algorithmic techniques for "next-item" prediction and discuss the results of a recent comparative evaluation of different algorithm families. In the final part of the talk, we will then take a closer look at the e-commerce domain. Specifically, we will report results from an in-depth study, which explored practical questions regarding the importance of short-term user intents, the use of recommendations as reminders, the role of community trends, and the recommendation of items that are on sale.

2 Sequence-Aware Recommender Systems

In [12], session-based recommendation is considered one main computational task of what is called *sequence-aware* recommender systems. Differently from traditional setups, the input to a sequence-aware recommendation problem is not a matrix of user-item preference statements, but a sequential log of past user interactions. Such logs, which are typically collected by today's e-commerce sites, can contain user interactions of various types such as item view events, purchases, or add-to-cart events.

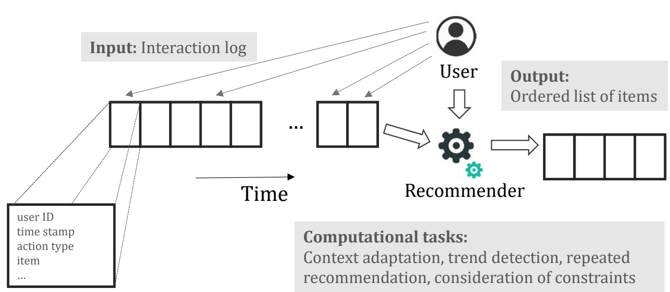

Fig. 1. Overview of the Sequence-Aware Recommendation Problem, adapted from [12].

Given such a log, various computational tasks can be defined. The most well-researched task in the literature is termed "context adaptation" in [12], where the goal is to create recommendations that suit the user's assumed short-term intents or contextual situation. Here, we can further discriminate between session-*based* and session-*aware* recommendation. In session-based scenarios, only the last few user interactions are known; in session-aware settings, in contrast, also past sessions of the current user might be available.

The sequential logs of sequence-aware recommender systems can however also be used for other types of computations, including the repeated recommendation of items, the detection of global trends in community, or the consideration of order constraints. These aspects are described in more detail in [12].

3 Session-based Recommendation

3.1 Algorithmic Approaches

A variety of algorithmic approaches have been proposed over the years for session-based recommendation scenarios. The conceptually most simple techniques rely on the detection of co-occurrence patterns in the recorded data. Recommendations of the form "Customers who bought ... also bought", as a simple form of session-based recommendation, can, for example, be determined by computing pairwise item co-occurrences or association rules of size two [1]. This concept can be extended to co-occurrence patterns that consider also the order of the events, e.g., in terms of simple Markov Chains or Sequential Patterns [11]. This latter approach falls into the category of *sequence learning* approaches [12], and a number of more advanced techniques based on Markov Decision Processes, Reinforcement Learning, and Recurrent Neural Networks were proposed in the literature [3, 13, 15]. In addition, distributional embeddings were explored to model user sessions in different domains. Finally, different hybrid approaches were investigated recently, which, for example, combine latent factor models with sequential information [14].

3.2 Evaluation Aspects: Recent Insights

Differently from the matrix completion problem formulation, no standards exist yet in the community for the comparative evaluation of session-based recommendation approaches, despite the existence of some proposals [5]. As a result, researchers use a variety of evaluation protocols and baselines in their experiments, which makes it difficult to assess the true value of new methods.

In [6] and [10], recently, an in-depth comparison of a variety of techniques for session-based recommendation was made. The comparison, which was based on datasets from several domains, included both conceptually simple techniques as well as the most recent algorithms based on Recurrent Neural Networks. To some surprise, it turned out that in almost all configurations, simple methods, e.g., based on the nearest-neighbor principle [6], where able to outperform the more complex ones. This, as a result, means that there is substantial room for improvement for more advanced machine learning techniques for the given problem setting.

4 On Short-Term Intents, Reminders, Trends, and Discounts in E-Commerce

In many session-based and session-aware recommendation problems in practice, a number of additional considerations can be made which are barely addressed in the academic literature. In [7], an in-depth analysis of various practical aspects was presented based on a large e-commerce dataset from the fashion domain.

The Role of Short-Term Intents. One first question relates to the *relative* importance of long-term preference models with respect to short-term user intents. The results presented, for example, in [5, 7] indicate that being able to estimate short-term intents is often much more important than further optimizing long-term preference models based, e.g., on matrix factorization techniques. One main challenge therefore lies in the proper estimation of the visitor's immediate shopping goal based only on a small set of interactions.

Recommendations as Reminders. While recommender systems in practice are often designed to also (repeatedly) recommend items that the user has inspected before, little research on the use of recommendations as reminders and navigation shortcuts exists so far. Recent research results however show that including reminders can have significant business value.

Trends and Discounts. A deeper analysis of a real-world dataset from the fashion domain in [7] furthermore reveals that recommending items that were recently popular, e.g., during the last day, is highly effective. At the same time, recommending items that are currently on sale leads to high click-to-purchase conversion, at least in the examined domain.

Learning Recommendations Success Factors from Log Data. A specific characteristic of the e-commerce dataset used in [7] is that it contains a detailed log of the items that were recommended to users along with information about clicks on such recommendations and subsequent purchases. Based on these logs, it is not only possible to analyze under which circumstances a recommendation was successful. We can also build predictive models based on these learned features, which at the end lead to more effective recommendation algorithms.

4.1 Challenges

Despite recent progress in the field, a variety of challenges remain to be further explored. Besides the development of more sophisticated algorithms for the next-item prediction problem, the open challenges, for example, include better mechanisms for combining long-term preference models with short-term user intents and to detect interest drifts. Furthermore, techniques can also be envisioned that are able to detect interest changes at the micro-level, i.e., during an individual session. In particular for the first few events in a new session, alternative approaches are needed to reliably estimate the user's short-term intent, based, e.g., on contextual information, global trends, meta-data, automatically extracted content-features, or from sensor information.

From a research perspective, the development of agreed-upon evaluation protocols and metrics are desirable, and more research is required to understand in which situation certain algorithms are advantageous. In addition, more user-oriented evaluations, as done in [9] for the music domain, are needed to better understand the utility of recommenders in different application scenarios.

From a more practical perspective, session-based recommendation can serve different purposes, e.g., they can be designed to either show alternatives options or complementary items. To be able to better assess the utility of the recommendations made by an algorithm for different stakeholders, purpose-oriented [4] and multi-metric evaluation approaches are required that go beyond the prediction of the next hidden item in offline experiments based on historical data.

References

1. Agrawal, R., Imieliński, T., Swami, A.: Mining association rules between sets of items in large databases. In: SIGMOD '93. pp. 207–216 (1993)
2. Ben-Shimon, D., Tsikinovsky, A., Friedmann, M., Shapira, B., Rokach, L., Hoerle, J.: Recsys challenge 2015 and the yoochoose dataset. In: ACM RecSys '15. pp. 357–358 (2015)
3. Hidasi, B., Karatzoglou, A., Baltrunas, L., Tikk, D.: Session-based recommendations with recurrent neural networks. In: ICLR '16 (2016)
4. Jannach, D., Adomavicius, G.: Recommendations with a purpose. In: RecSys '16. pp. 7–10 (2016)
5. Jannach, D., Lerche, L., Jugovac, M.: Adaptation and evaluation of recommendations for short-term shopping goals. In: RecSys '15. pp. 211–218 (2015)
6. Jannach, D., Ludewig, M.: When recurrent neural networks meet the neighborhood for session-based recommendation. In: RecSys '17. pp. 306–310 (2017)
7. Jannach, D., Ludewig, M., Lerche, L.: Session-based item recommendation in e-commerce: On short-term intents, reminders, trends, and discounts. User-Modeling and User-Adapted Interaction **27**(3–5), 351–392 (2017)
8. Jannach, D., Resnick, P., Tuzhilin, A., Zanker, M.: Recommender systems - beyond matrix completion. Communications of the ACM **59**(11), 94–102 (2016)
9. Kamehkhosh, I., Jannach, D.: User perception of next-track music recommendations. In: UMAP 2017 (2017)
10. Ludewig, M., Jannach, D.: Evaluation of session-based recommenation algorithms (2018), https://arxiv.org/abs/1803.09587
11. Mobasher, B., Dai, H., Luo, T., Nakagawa, M.: Using sequential and non-sequential patterns in predictive web usage mining tasks. In: ICDM 2003. pp. 669–672 (2002)
12. Quadrana, M., Cremonesi, P., Jannach, D.: Sequence-aware recommender systems. ACM Computing Surveys (2018)
13. Quadrana, M., Karatzoglou, A., Hidasi, B., Cremonesi, P.: Personalizing session-based recommendations with hierarchical recurrent neural networks. In: RecSys '17 (2017)
14. Rendle, S., Freudenthaler, C., Schmidt-Thieme, L.: Factorizing personalized markov chains for next-basket recommendation. In: WWW '10. pp. 811–820 (2010)
15. Shani, G., Heckerman, D., Brafman, R.I.: An MDP-based recommender system. J. Mach. Learn. Res. **6**, 1265–1295 (2005)