Exploiting multiple action types in recommender systems

Tural Gurbanov and Francesco Ricci

Free University of Bozen-Bolzano
tgurbanov@unibz.it

Abstract. Implicit feedback recommender systems provide personalized suggestions for items that are predicted to be of interest to the user, by collecting online users' activity and inferring from it users' preferences. While the state-of-the-art models are built by employing observations of user actions of one single type, the usage of multiple action types enables to observe more information and to build more effective recommenders. This paper presents an ongoing research in the field of implicit feedback recommender systems employing information about multiple action types. It highlights the peculiarities of existing models and identifies the research questions that should be solved to build solutions that can exploit multiple action types.

1 Introduction

Recommender systems are software tools and techniques providing personalized suggestions for items that are predicted to be of interest to the user [11]. The user may accept these suggestions or not and may provide, immediately or at a later stage, explicit or implicit feedback. Unlike explicit feedback, where a user's preference is provided explicitly (e.g., a rating), implicit feedback assumes that the system infers the user' opinion based on the user' actions (clicks, purchases, views, etc.). The advantage of implicit feedback is that there is no need for users to explicitly express their preferences (no cognitive load on the user) and that potentially every user interaction with the system can be considered as a feedback [1]. However, in comparison to explicit user feedback, which can be positive or negative, implicit feedback has no explicit valence [4] and only indirectly signals users' preferences or opinions [7]. Though high-quality explicit feedback is considered to be the most informative type of source data [4], in many real-world scenarios it can be difficult to obtain or unavailable (e.g., news portals). For this reason, more and more recommenders are built by leveraging implicit feedback data only.

To predict users' preferences by employing implicit feedback datasets the majority of existing models consider only actions of a single type. For example, Lin et al. [6] recommend crowd-sourcing tasks by using the number of times a user completed a given task. However, the interactions between a user and an item are very rarely limited to a single type of actions. and it has been recently shown that it is possible to improve a target action prediction (e.g, video views) by leveraging information about multiple action types (e.g., open details, and show later) [12, 3]. Although some works have been done in this direction, a general model has not been created yet.

In summary, this paper describes an ongoing research in the field of recommender systems employing implicit feedback in the form of a log of multiple action types. The primary motivation behind the research is the possibility to create more effective recommender systems that will not require users to specify their preferences explicitly but it will employ massive and presumably information-rich data generated by users' activity.

2 Multiple action types models

Though single action type models provide effective results, the availability of additional information (e.g., explicit feedback) and multiple action types makes it possible to improve action prediction accuracy [8, 10]. The collection of multiple action types has the potential to elicit more information about a user behaviour and to build more comprehensive user models. The resulting models can better describe users' preferences and allow creating better recommender systems.

In [3] we showed that action types correlated with a target action type (no matter positively or negatively) can be jointly employed to predict a target action. For example, analyzing a data set provided by a video streaming company we noted that if a user opened a video details page but did not watch the video then there is a high probability that the user will not watch the video in the future. Thus, if a user has not watched a video yet, then the observation of video details opened action decreases the confidence on the video watch action observation in the future.

The proposed general action prediction model considers a dataset of $d+1$ types of users' actions, where the variable a^j_{ui} counts or measures the observations of the action of type j performed by the user u on the item i. Actions of type $j = 0$ are considered to be the target data to predict and non-target action types $j \in \{1, \ldots, d\}$ are instead only contributing to the prediction. To predict whether u will act on i, with an action of the target type, the *indicator* and *confidence* values are introduced. The indicator value p_{ui} shows whether the user u has either performed the target action on the item i ($p_{ui} = 1$) or not ($p_{ui} = 0$), while c_{ui} indicates the confidence in the value assigned to p_{ui}. The indication function of the target action p_{ui} and the confidence c_{ui} of the target action are functions of all the observations:

$$p_{ui} = p(a^0_{ui}, a^1_{ui}, \ldots, a^d_{ui}) \qquad c_{ui} = c(a^0_{ui}, a^1_{ui}, \ldots, a^d_{ui})$$

The model generates predictions $\hat{p}_{ui}$ for the items i, such that $p_{ui} = 0$. The prediction is computed by using matrix factorization [5]: each user u and item i is associated with an f-dimensional factors vector $x_u \in \mathbb{R}^f$ and $y_i \in \mathbb{R}^f$ respectively. The predicted value of the indicator function is computed by the dot product of these two vectors: $\hat{p}_{ui} = x_u^T y_i$. The vectors' parameters are learned by minimizing the following cost function:

$$\min_{x_*, y_*} \sum_{u,i} c_{ui}(p_{ui} - x_u^T y_i)^2 + \lambda(\sum_u \|x_u\|^2 + \sum_i \|y_i\|^2)$$

The constant λ is used for tuning the regularization and avoiding model's overfitting [9]. Cost function minimization is achieved by alternating least squares optimization [4]. The empirical evaluation of the model was conducted on a large real-world dataset, and showed that using multiple actions is beneficial and can outperform a state-of-the-art implicit feedback model that uses only the target action data.

We note that the incorporation of multiple action types into the prediction requires the definition of specific indicator and confidence functions and the optimization of multiple meta parameters. In [3] these functions have been identified heuristically, after an exploratory analysis of the data. However, when several action types are present it may be difficult to manually explore which action types can help to predict a target action, especially, when the relation between predictive and target action types is subtle. To automatically identify the hidden relationships between actions of different types data mining, machine learning and time series analysis techniques can be applied.

To address the aforementioned problem, in [2] we introduced a hybrid recommender system for predicting user actions that employs Sequence Mining (SM) and Collaborative Filtering (CF) techniques. The system leverages observations of a range of actions, of different types, performed by the user on the target item and by other users on other items. The SM component considers the ordering of the actions and the time delays between them, while the CF component employs collaboration between multiple items and users. The limitation of the proposed hybrid model is that it does not exploit the impact of an action performed on an item, on actions performed on another item by the same user.

3 Problem issues

The usage of information about multiple action types in recommender systems is challenging because of the general limitations of implicit feedback [4], and the complexity of identifying which action types are relevant for the prediction of a target action type, and exactly determining the correlations between action of different types. In fact, the following research questions should be answered:

1. How to generate relevant recommendations by exploiting information about users' actions?
   Since the actions passively track users' behaviour, it is difficult to derive from them the users' real preferences and opinions.
2. How to discover automatically dependency relations between action types?
   The exploration of methods that automatically discover required action types and correlation between them is important and can improve recommendation results.

3. How to solicit actions that actually reveal unknown preference information? The recommender is not only a channel for delivering recommendations. In fact, recommendations generate reactions and tailor-made recommendations can be used to acquire the most informative feedback.

References

1. Claypool, M., Le, P., Wased, M., Brown, D.: Implicit interest indicators. In: Proceedings of the 6th International Conference on Intelligent User Interfaces. pp. 33–40. IUI '01, ACM, New York, NY, USA (2001), http://doi.acm.org/10.1145/359784.359836
2. Gurbanov, T., Ricci, F.: Action prediction models for recommender systems based on collaborative filtering and sequence mining hybridization. In: Proceedings of the 32nd ACM SIGAPP Symposium On Applied Computing. pp. 1655–1661. SAC'17 (2017)
3. Gurbanov, T., Ricci, F., Ploner, M.: Modeling and predicting user actions in recommender systems. In: Proceedings of the 2016 Conference on User Modeling Adaptation and Personalization. pp. 151–155. UMAP '16, ACM, New York, NY, USA (2016), http://doi.acm.org/10.1145/2930238.2930284
4. Hu, Y., Koren, Y., Volinsky, C.: Collaborative filtering for implicit feedback datasets. In: Proceedings of the 2008 Eighth IEEE International Conference on Data Mining. pp. 263–272. ICDM '08, IEEE Computer Society, Washington, DC, USA (2008), http://dx.doi.org/10.1109/ICDM.2008.22
5. Koren, Y., Bell, R., Volinsky, C.: Matrix factorization techniques for recommender systems. Computer 42(8), 30–37 (Aug 2009), http://dx.doi.org/10.1109/MC.2009.263
6. Lin, C.H., Kamar, E., Horvitz, E.: Signals in the silence: Models of implicit feedback in a recommendation system for crowdsourcing. In: Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence. pp. 908–914. AAAI'14, AAAI Press (2014), http://dl.acm.org/citation.cfm?id=2893873.2894015
7. Oard, D., Kim, J.: Implicit feedback for recommender systems. In: in Proceedings of the AAAI Workshop on Recommender Systems. pp. 81–83 (1998)
8. Parra, D., Amatriain, X.: Walk the talk: Analyzing the relation between implicit and explicit feedback for preference elicitation. In: Proceedings of the 19th International Conference on User Modeling, Adaption, and Personalization. pp. 255–268. UMAP'11, Springer-Verlag, Berlin, Heidelberg (2011), http://dl.acm.org/citation.cfm?id=2021855.2021878
9. Paterek, A.: Improving regularized singular value decomposition for collaborative filtering. Proceedings of KDD Cup and Workshop pp. 39–42 (2007)
10. Peska, L., Vojtas, P.: Using implicit preference relations to improve recommender systems. Journal on Data Semantics pp. 1–16 (2016), http://dx.doi.org/10.1007/s13740-016-0061-8
11. Ricci, F., Rokach, L., Shapira, B.: Recommender Systems Handbook, chap. Recommender Systems: Introduction and Challenges, pp. 1–34. Springer US, Boston, MA, 2nd edn. (2015), http://dx.doi.org/10.1007/978-1-4899-7637-6_1
12. Sun, Y., Yuan, N.J., Xie, X., McDonald, K., Zhang, R.: Collaborative nowcasting for contextual recommendation. In: Proceedings of the 25th International Conference on World Wide Web. pp. 1407–1418. WWW '16, International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, Switzerland (2016), http://dx.doi.org/10.1145/2872427.2874812