Continual Learning for Session-based Recommendation

Introduction

Due to new privacy regulations that prohibit building user preference models from historical user data, it is getting more and more important to utilize short-term dynamic preferences within an anonymous web browser session. Session-based recommendation systems (SRS, [1]) are therefore increasingly used in interactive online computing systems. The goal of SRS is to make recommendations based on user behavior observed in web browser sessions, and the task is to predict the users' next actions, such as clicks, based on previous actions in the current session.

Approaches proposed so far are all designed for a static setting where the user preference distribution and the set of items during the testing phase are assumed to be the same as in the training phase. Although such a static setting is standard for testing offline machine learning models, it does not reflect the dynamic nature of practical recommendation tasks.

To better address the dynamic nature of the SRS task, we study it from a realistic continual learning perspective. In this setting, new items and preferences appear continually during the testing phase. Common examples are news, forums and other social media, but also e-commerce. The challenge is to incorporate the new preference patterns incrementally while preserving static old ones that are still useful. This setting requires the recommender that is pre-trained with events observed during the last time period to be continually adapted to new items and user preferences in the following time period.

State-of-the-art neural approaches [1,2,3,4,5] have shown strong performance in standard static recommendation scenarios. Three challenges prevent them from being applied to continual SRS scenarios. First, their computation overhead prevents them from being frequently retrained or updated. Second, incorporating new items and preferences requires additional training and often reduces performance on preference patterns learned before. This is referred to as the catastrophic forgetting issue. Third, the small number of observations on new items and patterns cannot be captured well by these neural recommenders due to their statistical insignificance compared to other frequent patterns. This is referred to as the sample inefficiency issue.

A potential approach to work with continual recommendation scenarios is techniques from the **continual learning** community, which is also referred to as incremental learning or lifelong learning. Different techniques have been proposed to deal with the above three challenges, such as regularization/distillation methods [6,7,8] or memory-based rehearsal methods [9,10]. A complete paper list of methods in this popular area can be found at https://github.com/xialeiliu/Awesome-Incremental-Learning.

Your Job

In this project, you are going to adapt the state-of-the-art distillation and rehearsal techniques above to the continual recommendation task. You will be the forerunner to work on the continual learning issues for the recommendation task starting with formulating this problem, followed by building a benchmark model for this task.

Requirement

- **Familiar** with machine learning, especially deep learning concepts and math.
- Related project experience on recommendation system or deep learning is a plus.
- **Familiar** with Python and Pytorch (or Tensorflow).
- Be able to read and understand academic research papers.
- Self-motivated to on this topic.
- Please attach your CV and transcript.

Reference

[1] Balázs Hidasi, Alexandros Karatzoglou, Linas Baltrunas, and Domonkos Tikk. 2016. Session-based recommendations with recurrent neural networks. In Proceedings of 4th International Conference on Learning Representations (ICLR)
[2] Jing Li, Pengjie Ren, Zhumin Chen, Zhaochun Ren, Tao Lian, and Jun Ma. 2017. Neural attentive session-based recommendation. In the 26th ACM International Conference on Information and Knowledge Management (CIKM). ACM, 1419–1428.
[3] Qiao Liu, Yifu Zeng, Refuoe Mokhosi, and Haibin Zhang. 2018. STAMP: Short-term attention/memory priority model for session-based recommendation. In Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining. ACM, 1831–1839.
[4] Shu Wu, Yuyuan Tang, Yanqiao Zhu, Liang Wang, Xing Xie, and Tieniu Tan. 2019. Session-based Recommendation with Graph Neural Networks. In AAAI Conference on Artificial Intelligence (AAAI-19)
[5] Fei Sun, Jun Liu, Jian Wu, Changhua Pei, Xiao Lin, Wenwu Ou, and Peng Jiang. 2019. BERT4Rec: Sequential Recommendation with Bidirectional Encoder Representations from Transformer. In The 28th ACM International. Conference on Information and Knowledge Management (CIKM)
[6] James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska Barwinska, et al. Overcoming catastrophic forgetting in neural networks. Proceedings of the National Academy of Sciences, 114(13):3521–3526, 2017
[7] Zhizhong Li, and Derek Hoiem. Learning without forgetting. In European Conference on Computer Vision, pages 614– 629. Springer, 2016
[8] Hou, Saihui, et al. "Learning a Unified Classifier Incrementally via Rebalancing." CVPR 2019.
[9] Sylvestre-Alvise Rebuffi, Alexander Kolesnikov, Georg Sperl, and Christoph H. Lampert. icarl: Incremental classifier and representation learning. In The IEEE Conference on Computer Vision and Pattern Recognition (CVPR), July 2017.
[10] Francisco M. Castro, Manuel J. Marin-Jimenez, Nicolas Guil, Cordelia Schmid, and Karteek Alahari. End-to-end incremental learning. In The European Conference on Computer Vision (ECCV), September 2018.