Real-Time Recommendations in a Multi-Domain Environment

Emanuel Lacic
Know-Center & Graz University of Technology, Graz, Austria

Emanuel Lacić is a Software Engineer and Scientific Researcher in the Social Computing division at the Know-Center. Additionally, he is a PhD student at the Interactive Systems & Data Science Institute of Graz University of Technology. He has a M.Sc. and B.Sc. in Software Engineering and Information Systems from the Faculty of Electrical Engineering and Computing at the University of Zagreb. His PhD focus is on real-time recommender systems in large-scale settings. Apart from that, his research is utilized in several industry and European-funded projects such as Learning Layers, AFEL or MoreGrasp. His main research interests are in the fields of Recommender Systems, Deep Learning, Information Retrieval, Big Data, as well as Social and Complex Network Analysis. For more information please see `https://lacic.github.io/`

In the past decade, there has been a vast amount of research in the field of recommender systems. Most of that work focuses on developing novel approaches and improving accuracy. Thus, many well known methods are available, all having their unique strengths and weaknesses. These approaches are traditionally adapted and applied with the focus on a single domain model (e.g., marketplace, hotel, news, etc.). However, to support a diverse set of domains is becoming an important issue for modern recommender systems. In most domains, the prediction task is usually viewed as a two-dimensional problem which one needs to solve (e.g., utilizing user-item interactions). But nowadays it is not enough to support multiple domains on the basis of only one common data feature. With the arrival of the big data era, recommender systems are expected to analyze a lot of data, to support various data types and to handle streams of new data (i.e., volume, variety and velocity defining the Big Data problem). In such large-scale settings, traditional recommender systems usually analyze the data offline and update the generated model in regular time intervals. However, in many domains, choices made by users depend on factors which are susceptible to change anytime and being able to capture a user's real-time interests is gaining momentum.

In this PhD dissertation, it is first investigated to what extent different data features (i.e., item, social or location) can be utilized and combined for real-time recommendation. The experimental results not only suggest that performance can be gained by utilizing different data sources, but also that combining them should result into more robust recommendations. This is especially the case when recommending on different levels of specialization (i.e., categories). Building up on these results, a general framework is proposed which can (1) process streaming data online while providing real-time recommendations, (2) support

a multi-domain environment and the corresponding data features, and (3) provide a scalable architecture to cope with increasing request loads. For that purpose, it was shown that there is a benefit in using search engines (e.g., Apache Solr) as a recommender system is able to process data updates in real-time and immediately consider these updates (i.e., user's real-time interest) in the recommendation process without the need for recalculations. Finally, to provide a scalable and customizable architecture suited for multiple domains, a microservices architecture has been adopted. This way, different domains can run (and scale) the framework in isolated environments. The domain specific data features and recommendation approaches can be dynamically customized using a dedicated microservice which synchronizes the change to all domain-relevant nodes.