Recent Advances in Recommender Systems: Sets, Local Models, Coverage, and Errors

George Karypis

Department of Computer Science & Engineering,
University of Minnesota
Minneapolis, MN 55455, USA
karypis@umn.edu

ABSTRACT

Recommender systems are designed to identify the items that a user will like or find useful based on the user's prior preferences and activities. These systems have become ubiquitous and are an essential tool for information filtering and (e-)commerce. Over the years, collaborative filtering, which derive these recommendations by leveraging past activities of groups of users, has emerged as the most prominent approach for solving this problem. This talk will present some of our recent work towards improving the performance of collaborative filtering-based recommender systems and understanding some of their fundamental limitations and characteristics. It will start by analyzing how the ratings that users provide to a set of items relate to their ratings of the set's individual items and, using these insights, will present rating prediction approaches that utilize distant supervision. It will then discuss extensions to approaches based on sparse linear and latent factor models that postulate that users' preferences are a combination of global and local preferences, which are shown to lead to better user modeling and as such improved prediction performance. Finally, the talk will conclude by discussing what can be accurately predicted by latent factor approaches and by analyzing the estimation error of sparse linear and latent factor models and how its characteristics impacts the performance of top N recommendation algorithms.

CCS Concepts/ACM Classifiers

- Information systems~Collaborative filtering
- Information systems~Personalization
- Computing methodologies~Multi-task learning

Author Keywords

Distant supervision; Latent factor models; Item-based models

BIOGRAPHY

George Karypis is a Distinguished McKnight University Professor and an ADC Chair of Digital Technology at the Department of Computer Science & Engineering at the University of Minnesota, Twin Cities. His research interests span the areas of data mining, high performance computing, information retrieval, collaborative filtering, bioinformatics, cheminformatics, and scientific computing. His research has resulted in the development of software libraries for serial and parallel graph partitioning (METIS and ParMETIS), hypergraph partitioning (hMETIS), for parallel Cholesky factorization (PSPASES), for collaborative filtering-based recommendation algorithms (SUGGEST), clustering high dimensional datasets (CLUTO), finding frequent patterns in diverse datasets (PAFI), and for protein secondary structure prediction (YASSPP). He has coauthored over 280 papers on these topics and two books ("Introduction to *Protein Structure Prediction: Methods and Algorithms*" (Wiley, 2010) and "Introduction to Parallel Computing" (Publ. Addison Wesley, 2003, 2nd edition)). In addition, he is serving on the program committees of many conferences and workshops on these topics, and on the editorial boards of the IEEE Transactions on Knowledge and Data Engineering, ACM Transactions on Knowledge Discovery from Data, Data Mining and Knowledge Discovery, Social Network Analysis and Data Mining Journal, International Journal of Data Mining and Bioinformatics, the journal on Current Proteomics, Advances in Bioinformatics, and Biomedicine and Biotechnology. He is a Fellow of the IEEE.

This paper is published under the Creative Commons Attribution-NonCommercial-NoDerivs 4.0 International (CC BY 4.0) license. Authors reserve their rights to disseminate the work on their personal and corporate Web sites with the appropriate attribution.

WWW '18 Companion April 23-27, 2018, Lyon, France.

© 2018 IW3C2 (International World Wide Web Conference Committee), published under Creative Commons CC BY 4.0 License.

ACM ISBN 978-1-4503-5640-4/18/04.

DOI: <https://doi.org/10.1145/3184558.3191588>