

Large-Scale Social Recommender Systems: Challenges and Opportunities

Mitul Tiwari
LinkedIn
mtiwari@linkedin.com

ABSTRACT

Online social networks have become very important for networking, communication, sharing, and content discovery. Recommender systems play a significant role on any online social network for engaging members, recruiting new members, and recommending other members to connect with. This talk presents challenges in recommender systems, graph analysis, social stream relevance and virality on a large-scale social networks such as LinkedIn, the largest professional network with more than 200M members.

First, social recommender systems for recommending jobs, groups, companies to follow, other members to connect with, are very important part of a professional network like LinkedIn [1, 6, 7, 9]. Each one of these entity recommender systems present novel challenges to use social and member generated data. Second, various problems, such as, link prediction, visualizing connection network, finding the strength of each connection, and the best path among members, require large-scale social graph analysis, and present unique research opportunities [2, 5]. Third, social stream relevance and capturing virality in social products are crucial for engaging users on any online social network [4]. Final, systems challenges must be addressed in scaling recommender systems on a large-scale social networks [3, 8, 10]. This talk presents challenges and interesting problems in large-scale social recommender systems, and describes some of the solutions.

Categories and Subject Descriptors

H.2.8 [Database Applications]: Data Mining

Keywords

Recommender Systems, Social Networks, Link Prediction, Relevance

Bio

Mitul Tiwari is a Staff Research Engineer and Tech Lead at Search, Network, and Analytics group at LinkedIn. He is working on large-scale social recommender systems such as “People You May Know” and “Related Searches”. His interests include recommender systems, social network analysis, large-scale data mining, machine learning, and distributed systems. Previously, he worked at Kosmix (now Walmart Labs) as a Lead Member of Technical Staff, where he worked on web-scale document and query categorization, and its application to vertical search, topic pages, and tweets classification. He completed his PhD in Computer Science from the University of Texas at Austin. Earlier he received his under graduation from Indian Institute of Technology, Bombay. He has co-authored more than a dozen publications in conferences such as WWW, SIGIR, CIKM, and SPAA. Find more details here: <http://www.mitultiwari.net>.

References

- [1] M. S. Amin, B. Yan, S. Sriram, A. Bhasin, and C. Posse. Social referral: leveraging network connections to deliver recommendations. In *Proceedings of the 6th ACM Conference on Recommender Systems*, 2012.
- [2] S. Budalakoti and R. Bekkerman. Bimodal invitation-navigation fair bets model for authority identification in a social network. In *Proceedings of the 21st International World Wide Web Conference*, 2012.
- [3] K. Goodhope, J. Koshy, J. Kreps, N. Narkhede, R. Park, J. Rao, and V. Y. Ye. Building linkedin’s real-time activity data pipeline. *IEEE Data Engineering Bulletin*, 35(2):33–45, 2012.
- [4] L. Hong, R. Bekkerman, J. Adler, and B. D. Davison. Learning to rank social update streams. In *Proceedings of the 35th International ACM SIGIR Conference*, 2012.
- [5] C.-J. Hsieh, M. Tiwari, D. Agarwal, X. L. Huang, and S. Shah. Organizational overlap on social networks and its applications. In *Proceedings of the 22nd International World Wide Web Conference*, 2013.
- [6] A. Reda, Y. Park, M. Tiwari, C. Posse, and S. Shah. Metaphor: a system for related search recommendations. In *Proceedings of the 21st ACM International Conference on Information and Knowledge Management*, 2012.
- [7] M. Rodríguez, C. Posse, and E. Zhang. Multiple objective optimization in recommender systems. In *Proceedings of the 6th ACM Conference on Recommender Systems*, 2012.
- [8] R. Sumbaly, J. Kreps, L. Gao, A. Feinberg, C. Soman, and S. Shah. Serving large-scale batch computed data with project voldemort. In *Proceedings of the 10th USENIX conference on File and Storage Technologies*, 2012.
- [9] J. Wang, Y. Zhang, C. Posse, and A. Bhasin. Is it time for a career switch. In *Proceedings of the 22nd International World Wide Web Conference*, 2013.
- [10] L. Wu, R. Sumbaly, C. Riccomini, G. Koo, H. J. Kim, J. Kreps, and S. Shah. Avatara: Olap for web-scale analytics products. *Proceedings of VLDB Endowment*, 5(12):1874–1877, 2012.