Personalization in Social Networks: Modeling the Underlying Social Processes

Amit Sharma, Dept. of Computer Science, Cornell University

Abstract

What is the connection between our online social ties and preferences in music, movies and other domains? Theories from sociology such as homophily and influence and recent empirical studies on the web suggest a *locality* in preference, that is, people close to each other socially are likely to have similar preferences. However, we find that the locality in preference varies from person to person and community to community---which is perhaps the reason why social recommendation has led to mixed results. A fruitful direction of research can be to understand the underlying social processes that likely affect user preferences in online social networks. A deeper understanding of these social processes around user preference can help in creating better models of preference and lead to more effective personalization strategies.

Introduction

In recent years, online platforms have emerged which thrive on the activity of their users who share some social connections. For example, users share updates with their followers on Twitter, rate and discover music on Grooveshark or Spotify, edit articles on Wikipedia and post text, pictures and comments on social networks like Facebook and Google+. In each of these domains, people interact and form social connections with other users. At the same time they also express their *preferences* towards items through their rating, sharing, editing, or consuming activities. Availability of this additional social context has opened up new avenues for personalization---closeness in a social network likely implies similar preferences (locality of preference) [11, 12, 16], an ability to influence those preferences [9] and greater chances of group consumption [18].

However, research on social recommendation by either augmenting collaborative filtering with social data [2, 5, 7] or using the first-degree connections for recommendation [3, 4, 6] has reported varying results on different networks and often the gains of using additional social information are only slight [1, 8]. One of the reasons could be that little is understood about the connection between people's explicit social connections and their interests in a particular domain---how do a user's social contacts influence his/her interests, which social relationships matter, and for what domains of interest?

For offline networks, social scientists have answered these questions by building theories of human behavior around preferences. These theories suggest that people are influenced by the preferences of their social contacts [9, 10], form social connections based on common preferences [11, 12], actively compare their preferences to others' preferences [13], conform their preferences to the larger consensus [14], or curate their preferences as a form of identity [15]. Understanding how these social processes operate and affect preferences in online social networks can be a fruitful direction of research that marries computational models of diffusion and network contagion with qualitative and behavioral analysis of people's activities. A deeper understanding of the connection between social processes and user preference would also lead to more effective models of personalization.

Towards effective personalization in social networks

We are conducting a series of studies towards understanding the role of social processes in preference evolution by using both small-scale experiments and large-scale data mining. Small-scale behavioral experiments give insights on how people think about, rate and make decisions about items in social contexts, while large-scale data analysis can shed light on how user preferences evolve and how item preferences are distributed in social networks. Next, we briefly describe some of our efforts in this direction.

To compare and evaluate the effect of social processes such as influence and conformity on users' rating of items, we conducted a user study in which we showed different types of additional social information (or *explanations*) for a musical artist recommendation, such as "Amit likes this" or "7 of your friends like this" [17]. The results using 237 Facebook users show that influence from close friends exceeds the conformity effects of showing the number of friends who like an artist. Further, the relative effect of social explanations varied with different people and even though we showed minimum information about the musical artists, people still considered the artist details for determining their rating. We formalized these notions by proposing a generative mixture model for producing a rating that combines a user's own expectation of liking an item with the effect of additional social information. These results have both practical and theoretical implications. On the practical side, our generative model can be used as a basis for personalizing explanation strategies. At the same time, it offers a window into users' decision-making processes around recommendations.

Analysis of datasets from Facebook and Twitter involving people's Likes in music, movies and usage of hashtags reveals more about the connection between social ties and people's preferences [16]. These datasets were collected as a part of studies conducted at Cornell [17, 19] and Stanford [20] universities and contain ego networks for a sample of Facebook and Twitter users. For all three domains, we found evidence of preference locality: on average, friends of a user do have more preferences in common than strangers. This was also reflected in recommendation algorithms for predicting a user's future likes or usage of hashtags; using only friends' past preferences gave comparable performance than using data from the entire available dataset. Considering that the data from all users is orders of magnitude larger than that from just friends, these results are of practical import for designers of personalization systems in social contexts. In addition, we found that the extent of preference locality varies with different domains which merits further explanation.

Integrating small-scale with large-scale

Insights from small-scale studies may inform our explanation and models of the effects seen in larger datasets. For instance, our model on how people rate recommendations from their social connections could be used to simulate the generation of people's Likes and explain the locality effects that we observed in datasets from Facebook and Twitter. Similarly, we are presently conducting a study to learn about the sharing behavior of users; insights about the processes by which people share could explain observations of item diffusion in large network datasets. Computationally accounting for social processes through a combination of small-scale behavioral exploration and large-scale data analysis can be a new paradigm for designing personalization models for individuals and groups. This, however, requires a cross-disciplinary and multi-method approach and we hope researchers will join us in advancing personalization in social contexts through a better understanding of the underlying social processes.

References

[1] G. Groh, S. Birnkammerer, and V. Kllhofer, “Social recommender systems,” in Recommender Systems for the Social Web, Intelligent Sys. Ref. Library, Springer Berlin Heidelberg, 2012.
[2] H. Ma, D. Zhou, C. Liu, M. R. Lyu, and I. King, “Recommender systems with social regularization,” in Proc. WSDM, 2011.
[3] I. Guy, N. Zwerdling, D. Carmel, I. Ronen, E. Uziel, S. Yogev, and S. Ofek-Koifman, “Personalized recommendation of social software items based on social relations,” in Proc. RecSys, 2009.
[4] J. Chen, R. Nairn, L. Nelson, M. Bernstein, and E. Chi, “Short and tweet: Experiments on recommending content from information streams,” in Proc. CHI, 2010.
[5] H. Ma, I. King, and M. R. Lyu, “Learning to recommend with social trust ensemble,” in Proc. SIGIR, 2009.
[6] A. Sharma and D. Cosley, “Network-centric recommendation: Personalization with and in social networks,” in Proc. IEEE SocialCom, 2011.
[7] I. Konstas, V. Stathopoulos, and J. M. Jose, “On social networks and collaborative recommendation,” in Proc. SIGIR, 2009.
[8] O. Arazy, N. Kumar, and B. Shapira, “A theory-driven design framework for social recommender systems,” in Journal of the Assoc. for Info. Sys., 2010.
[9] R. Cialdini, Influence: Science and practice, vol. 4. Allyn and Bacon Boston, MA, 2001.
[10] B. Latane, “The psychology of social impact.,” American psychologist, vol. 36, no. 4, p. 343, 1981.
[11] M. McPherson, L. Smith-Lovin, and J. M. Cook, “Birds of a feather: Homophily in social networks,” Annual Review of Sociology, vol. 27, no. 1, pp. 415–444, 2001.
[12] G. Kossinets and D. J. Watts, “Origins of homophily in an evolving social network,” American Journal of Sociology, vol. 115, 2009.
[13] J. Suls, R. Martin, and L. Wheeler, “Social comparison: Why, with whom, and with what effect?,” vol. 11, no. 5, pp. 159–163, 2002.
[14] R. B. Cialdini and N. J. Goldstein, “Social influence: Compliance and conformity,” Annu. Rev. Psychol., vol. 55, pp. 591–621, 2004.
[15] E. Pearson, “All the world wide webs a stage: The performance of identity in online social networks,” First Monday, vol. 14, no. 3, 2009.
[16] A. Sharma, M. Gemici, and D. Cosley, “Friends, strangers, and the value of ego networks for recommendation,” in Proc. ICWSM, 2013.
[17] A. Sharma and D. Cosley, “Do social explanations work? Studying and modeling the effects of social explanations in recommender systems,” in Proc. WWW, 2013.
[18] M. Gartrell, X. Xing, Q. Lv, A. Beach, R. Han, S. Mishra, and K. Seada, “Enhancing group recommendation by incorporating social relationship interactions”, in Proc. GROUP, 2010.
[19] A. Sharma and D. Cosley, “Network-centric recommendation: Personalization with and in social networks,” in Proc. IEEE SocialCom, 2011.
[20] J. J. McAuley and J. Leskovec, “Learning to discover social circles in ego networks,” in NIPS, 2012.