
GroupLink: Group Event Recommendations Using Personal Digital Traces

Honghao Wei

Tsinghua University
Beijing, 100084, China
weihh12@mails.tsinghua.edu.cn

Longqi Yang

Cornell Tech
New York, NY 10011, USA
ylongqi@cs.cornell.edu

Cheng-Kang Hsieh

UCLA CSD
Los Angeles, CA 90095, USA
changun@cs.ucla.edu

Deborah Estrin

Cornell Tech
New York, NY 10011, USA
destrin@cs.cornell.edu

Abstract

We present GroupLink, a group event recommendation system that suggests events to promote group members' face-to-face interactions in non-work settings. GroupLink addresses the challenge of finding events that appeal to a collection of individuals with diverse interests. The system mitigates the difficulties of identifying and tracking group members' preferences through analyzing and learning preferences from individual member's personal digital traces, including social media, email, and online streaming histories. It is a web-based service that can be easily accessed using diverse devices, and is particularly suited for enhancing interpersonal interactions in CSCW.

Author Keywords

Recommender Systems, Personal Digital Traces, Immersive Recommendation, Group Event Recommendation

ACM Classification Keywords

H.4.m [Information System Applications]: Miscellaneous

Introduction

Interpersonal interactions, effective collaboration, knowledge sharing, and trust among group members are critical for a group's success [1, 8]. One recommended way to encourage such interactions is through organizing group events [4]; from attending local meetups and community

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the Owner/Author.

Copyright is held by the owner/author(s).

CSCW '16 Companion, February 27 – March 02, 2016, San Francisco, CA, USA
ACM 978-1-4503-3950-6/16/02.

<http://dx.doi.org/10.1145/2818052.2874338>

works, to office outings and creative activities. These events allow group members to have face-to-face communications outside of office settings. Prior studies have shown that engaging in these non-work-related activities is associated with increased creative problem-solving at work, as well as greater support of coworkers [5], and the face-to-face communications can enhance the connections between group members and their sense of belongingness [1].

However, two significant challenges hinder a group to effectively plan such events. First, finding events that suit every member's interests is challenging, especially for the groups consisting of members with diverse background and interests [2]. A group event that fits only a few people's interests is clearly not ideal; but an event that appeals to the majority of people but is disliked by a few members is also undesirable as it may lead to the creation of "subgroups" or the "lone wolf" phenomenon that is detrimental to the group performance [3]. Also, a generic event that moderately appeals to everyone but is not really compelling for anyone is also not an desirable option.

The other challenge is that, due to the dynamic nature of group membership, an event organizer would have difficulty maintaining a thorough understanding of each member's preferences, particularly for new members of the group. One option is to use surveys to elicit the members' interests. However, besides the additional effort required for each group member, such an approach still requires the organizer to come up with events given everyone's preferences, and to keep up with the latest interests of all group members and incorporate the latest event options.

We propose **GroupLink**, an event recommender system, to address these challenges. Given a small-to-medium-sized group (3-8 people), GroupLink recommends the events that can best connect every group member. More specifically,

we extend the traditional *least-misery* strategy for the movie recommendation [7] to the scenario for event recommendations, called *best-minimum-connected* strategy, where the utility of a collection of recommended events is defined as the minimum number of opportunities those events will create for a member to interact with each other member. We adopt a greedy algorithm to optimize the expected *minimum-connectedness* of the group in order to better enhance the group interactions and prevent the subgroup and lone-wolf phenomena.

To address the difficulty of tracking each member's preferences, we rely on a new recommendation paradigm, called *Immersive Recommendation*, to allow a group of people, even if they just meet each other, to receive instant tailored group recommendations. Proposed in [6], immersive recommendation learns users' interests through their personal digital traces, such as email, social media, and online streaming watch histories; thereby avoiding the tedious user surveys or lengthy preference learning processes required by most recommendation systems. Immersive recommendation has been shown to achieve better than state-of-the-art recommendation accuracy for individual user local event recommendations (see [6] and <https://bit.ly/group-link> for a live demo). In GroupLink, we extend this recommendation model to group recommendation for small-to-medium-sized groups and study its effectiveness and the potential privacy concerns.

Research Questions

Through GroupLink, we intend to answer two research questions. First, how effective the proposed recommendation strategy is in terms of enhancing interactions and trust between group members (compared to the prior recommendation strategies such as least-misery or mean-preference) in this particular application context. Second, how feasible it

is to apply immersive recommendation to group recommendation to enable instant recommendations and its potential privacy concerns. The original immersive recommendation model assumes that recommendations are only visible to the users themselves. How much will users be willing to trade privacy for utility in the group recommendation setting given that their preferences might be implicitly exposed to the other group members [7]. Moreover, how can we design the user interface to better balance the privacy and usability of the recommendations.



Figure 1: *GroupLink* System Prototype. (I) User login page; (II) Events recommendation page; (III) Action page;

System Prototype

We prototype the proposed system, *GroupLink*, as a responsive web app, which can be easily accessed using diverse devices and platforms. The system consists of four phases: **Onboarding**, **Group Formation**, **Recommendation**, and **Action**, which will be detailed next.

Onboarding: During the onboarding process, as Fig.1 shows, each user can choose to opt-in at least one of their

digital traces, including Email, Twitter, Facebook, and YouTube watch history. These traces will be securely pulled from the third-party services with appropriate authorization enabled by OAuth protocol. Given these traces, we use the existing immersive recommendation engine developed in [6] to create a profile for each individual user.

Group Formation: Every user can create a group and invite others to join through email or QR codes, and a user can simultaneously belong to multiple groups they choose to join. After a user signs in, she will be prompted with a list of groups she belongs to and can choose one of them to see the recommendations tailored for that particular group.

Recommendation: The recommendations are made according to the expected preferences each member would have to each local event. The preferences are inferred from the individual users' profiles learned in the onboarding process (see [6] for more on individual preference learning). Assuming a user will attend an event if her expected preference is over a threshold, we use a greedy algorithm to search for a collection of events that optimize the *most-minimum-connected* metric mentioned previously. We will recommend events from *Meetup.com* that are in the San Francisco area as it is where CSCW will take place. As Fig.1 shows, for each event, we present basic information including *name*, *time* and *place*, along with a list of members who have joined the event.

Action: To further encourage the user participation, the system allows users to make in-channel communication and make comments to the recommended events. As Fig.1 shows, when the user clicks on any of the recommended events shown in the previous phase, the system will enter the **Action** phase where users can express the aspiration as to whether or not to join the event, provide more details about the events, and plan the logistics.

The current prototype is well suited for live demonstration over small-to-medium-sized groups with a wide range of diversities. In the future we will close the feedback loop by tracking whether opt-in users execute their expressed aspirations to attend meetups.

Conclusion

GroupLink is a group event recommender system, designed to address the challenges of finding events that suit different group members' interests, and increase the opportunities of face-to-face interactions among group members. It avoids tedious user surveys and lengthy learning processes by profiling users' interests from their personal digital traces, and generate recommendations that can better connect a group through a *best-minimum-connected* strategy. We prototype the proposed system as a responsive web app that is well suited for live demonstration in CSCW, where many people may attend as a group or form groups during the venue.

Acknowledgements

This research is partly funded by AOL-Program for Connected Experiences and further supported by the small data lab at Cornell Tech, which receives funding from UnitedHealth Group, Google, Pfizer, RWJF, NIH and NSF.

REFERENCES

1. Lisa C Abrams, Rob Cross, Eric Lesser, and Daniel Z Levin. 2003. Nurturing interpersonal trust in knowledge-sharing networks. *The Academy of Management Executive* 17, 4 (2003), 64–77.
2. Linas Baltrunas, Tadas Makcinskas, and Francesco Ricci. 2010. Group recommendations with rank aggregation and collaborative filtering. In *Proceedings of the fourth ACM conference on Recommender systems*. ACM, 119–126.
3. Tracy L Dumas, Katherine W Phillips, and Nancy P Rothbard. 2013. Getting closer at the company party: Integration experiences, racial dissimilarity, and workplace relationships. *Organization Science* 24, 5 (2013), 1377–1401.
4. Nurcan Karamolla Ensari and Norman Miller. 2006. The application of the personalization model in diversity management. *Group Processes & Intergroup Relations* 9, 4 (2006), 589–607.
5. Kevin J Eschleman, Jamie Madsen, Gene Alarcon, and Alex Barelka. 2014. Benefiting from creative activity: The positive relationships between creative activity, recovery experiences, and performance-related outcomes. *Journal of Occupational and Organizational Psychology* 87, 3 (2014), 579–598.
6. Cheng-Kang Hsieh, Longqi Yang, Honghao Wei, Mor Naaman, and Deborah Estrin. 2016. Immersive Recommendation: News and Event Recommendations Using Personal Digital Traces. In *Proceedings of the 25th ACM International Conference on World Wide Web*. ACM.
7. Mark O'connor, Dan Cosley, Joseph A Konstan, and John Riedl. 2001. PolyLens: a recommender system for groups of users. In *ECSCW 2001*. Springer, 199–218.
8. Jun-Gi Park and Jungwoo Lee. 2014. Knowledge sharing in information systems development projects: Explicating the role of dependence and trust. *International Journal of Project Management* 32, 1 (2014), 153–165.