Increasing Activity in Enterprise Online Communities Using Content Recommendation

IDO GUY, Ben Gurion University of the Negev and Yahoo Research
INBAL RONEN, IBM Research
ELAD KRAVI, Technion, Israel Institute of Technology
MAYA BARNEA, IBM Research

22

Although online communities have become popular both on the web and within enterprises, many of them often experience low levels of activity and engagement from their members. Previous studies identified the important role of community leaders in maintaining the health and vitality of their communities. One of their key means for doing so is by contributing relevant content to the community. In this paper, we study the effects of recommending social media content on enterprise community leaders. We conducted a large-scale user survey with four recommendation rounds, in which community leaders indicated their willingness to share social media items with their communities. They also had the option to instantly share these items. Recommendations were generated based on seven types of community interest profiles that were member-based, content-based, or hybrid. Our results attest that providing content recommendations to leaders can help uplift activity within their communities.

CCS Concepts: • **Human-centered computing** → **Social recommendation**; *User models*; *User studies*; *Collaborative and social computing*; *Social tagging*; *Social media*; Social networks; Social networking sites; • **Information systems** → *Enterprise applications*; *Personalization*; *Recommender systems*; • **Networks** → Online social networks

Additional Key Words and Phrases: Enterprise, engagement, group recommendation, online communities, social recommender systems, social business

ACM Reference Format:
Ido Guy, Inbal Ronen, Elad Kravi, and Maya Barnea. 2016. Increasing activity in enterprise online communities using content recommendation. ACM Trans. Comput.-Hum. Interact. 23, 4, Article 22 (August 2016), 28 pages.
DOI: http://dx.doi.org/10.1145/2910581

1. INTRODUCTION

Online communities have become very popular in recent years, both on the internet and in the enterprise. For example, Matthews et al. [2013] report the existence of 111,577 communities with 487,941 distinct members within IBM's social media environment, indicating that almost every employee is a member of at least one community. Organizations make use of communities to increase productivity and share expertise across global teams. Although many online communities started as forums for question and answering or idea sharing, they have begun to include additional social media tools, such as blogs, wikis, microblogs, and shared files. Such tools enable the community members to create and share content in diverse ways according to their needs [Matthews et al. 2014b].

Authors' addresses: I. Guy, Matam Park, Haifa 31905, Israel; email: idoguy@acm.org; I. Ronen and M. Barnea, Haifa University, Mt. Carmel, Haifa 34988, Israel; emails: {inbal, mayab}@il.ibm.com; E. Kravi, Technion Campus, 32000, Israel; email: ekravi@cs.technion.ac.il.

DOI: http://dx.doi.org/10.1145/2910581

In the enterprise, communities serve a unique role. Muller et al. [2012] highlight three "critical differences" between enterprise and web communities: (1) an enterprise provides a shared context in addition to the context of the community, which can contribute to a level of trust and common ground; (2) enterprise communities are typically business-focused, leading to different contents and styles of discussion; and (3) companies, which require authenticated access and use of real names, eliminate anonymity and provide greater transparency.

The potential benefits of online communities to the employees and the business are numerous and may include breaking down organizational and distance barriers to knowledge sharing and collaboration, improved skills and ability to execute and retain staff, improved sales, improved speed of execution, facilitation of team work, and enhanced innovation processes [Matthews et al. 2014b; Matthews et al. 2013].

Despite their proliferation and potential benefits, many enterprise online communities face a challenge of getting their members to participate and contribute and often fail to strive [Ehrlich et al. 2014; Xu et al. 2013]. For example, Zhu et al. [2014] examined a set of 10,000 most-recently updated communities in a large organization and found that they had an average of only 10 activities during a period of 2.5 months. Low participation rates are likely due to the growing number of enterprise communities, as well as the fact that in a workplace, employees need to carefully choose where they invest their time; Leftheriotis and Giannakos [2014] and Ehrlich et al. [2014] stated, "*Even in an enterprise setting where online communities augment other business tools, it is challenging to have members make the connection between the level of community activity and their own work.*"

Various studies identified features characterizing healthy online communities and strategies of how to encourage and maintain lively communities [Bateman et al. 2011; Iriberry and Leroy 2009; Kim 2000]. Among others, it has been shown that high-quality and up-to-date content is a significant contributor to the success of a community. Its effects can include heightened member activity [Koh et al. 2007], which is commonly used as an indicator of a community's health and success [Ehrlich et al. 2014; Iriberry and Leroy 2009; Matthews et al. 2014a; Muller et al. 2012; Preece et al. 2004].

Community owners, also commonly referred to as moderators or leaders, face the challenge of keeping their communities alive and relevant. As many communities suffer from low member engagement, owners often have a hard time supporting the members' needs and expectations. Previous studies highlighted the impact of the community owner on ensuring the proper evolution of the community [Gray 2005; Koh et al. 2007]. Engaged owners, who take the role upon themselves, contribute significantly to a community's success [Ehrlich et al. 2012]. For example, owners can promote the success of the community by encouraging contributions and discussions and contributing content themselves [Butler et al. 2007; Kim 2000; Koh et al. 2007]. Supporting the owner with tools that help manage and boost the community is therefore essential for the facilitation of a valuable community, which takes the advantage of the potential benefits mentioned before [Matthews et al. 2013; Xu et al. 2013].

In this paper, we use a recommender system approach to aid owners with their task of sharing content with the community and ultimately get more members to participate. Specifically, we study the recommendation of social media content that enterprise community owners can share with their community. We examine several community interest profiles, which are based on members' interests, the content of the community, or a hybridization of both. Our recommender system is based on a method from a previous study that explored content recommendation to community owners by comparing owners' interest ratings for themselves versus for the whole community [Ronen et al. 2014]. Our work goes beyond interest ratings to examine the actual sharing action and its influence on the community.

To test the quality of the recommendations and their effect on community activity, we sent four rounds of recommendations to owners of active communities in an enterprise, over a period of six weeks. We received responses from 1,053 owners of 903 communities in the first round, down to 183 owners of 177 communities in the fourth round. Owners were presented with social media items such as blog posts, wiki pages, and bookmarks relevant to their community according to one of seven interest profiles that we experimented with. For each item, the owners were asked to indicate whether they would share it with their community. We also provided the ability to instantly share the item with the community by clicking a button in the survey. In our evaluation, we compared the impact of the different interest profiles on the motivation of owners to share the item within the community and on their actual sharing actions. We also examined how the characteristics of the recommended items, such as their application source, influenced recommendation effectiveness. In addition, we investigated the impact of the recommendations on the general activity level of the community and the number of contributing members over the eight weeks that followed the beginning of our survey.

Our results indicate that recommendations based on the community's content are especially effective in influencing owners' willingness to share them with the community and take real action. We also found that our survey had a significant positive effect on community activity during the eight weeks following its initiation. This effect was reflected both through the sheer volume of activity and the number of contributing members. Overall, our findings suggest that recommendations to community owners can help raise activity and engagement within the community.

2. RELATED WORK

Our related literature review is divided into three areas: (1) enterprise online communities, who are the target of our study; (2) recommendation to groups and communities, which relate to our recommendation approach; and (3) engagement enhancement in social media systems, since a key task in our research focuses on measuring increase in community activity and member participation.

2.1. Enterprise Online Communities

Online communities are increasingly being deployed in enterprises to grow productivity and share expertise [Matthews et al. 2013]. Early work on enterprise online communities presented case studies of small numbers of organizationally sponsored online communities [Millen et al. 2002; Wenger and Snyder 2002; Ebrahim et al. 2009] that provide a good summary. In the recent few years, with the growing popularity of both enterprise social media and enterprise online communities, various new studies have emerged. Muller et al. [2012] examined different types of enterprise communities and how they make the use of social media tools. The two most common types were communities of practice (a group of people with a common interest or practice) and teams (working on a shared goal for a particular customer, project, or business function). In addition, three other types were identified: technical support, idea labs (for brainstorming), and recreation communities (for leisure activities). Matthews et al. [2014b] analyzed and surveyed members of 128 workplace communities to examine the use of different social media tools and found that different communities exercise different combinations and practices. In a qualitative study, Ehrlich et al. [2014] identified "informal leaders" in enterprise online communities and found they were mostly motivated by wanting to help other members and gaining access to information resources.

Quite a few studies examined practices and measurements for community success and engagement [Bateman et al. 2011; Fugelstad et al. 2012; Kim 2000; Ma and Agarwal 2007]. Several highlighted the importance of high-quality content, which can lead to an increased activity of the members and plays a key role in the assessment of a community's health; Butler et al. [2007], Iriberri and Leroy [2009], and Koh et al. [2007] stated that useful content leads community members to increase their activity. In a field study of 77 virtual Korean communities, they found that content usefulness increased viewing and posting activity.

Leimeister et al. [2004] showed, using surveys, that high-quality and up-to-date content should be provided to maintain the success of a community. Zhu et al. [2014] found that communities that linked to content of other communities with overlapping topics had an increased level of activity. Community insights [Matthews et al. 2013] provided actionable analytics that helped community leaders foster healthy communities. One of the main actions proposed was the need to identify critical content by owners.

The critical role of owners in a community's success has been underlined in many studies [Bourhis et al. 2005; Gray 2005; Huffaker 2010]. Sangwan [2005] surveyed member satisfaction within a financial knowledge community and stated that the content generated by a community owner, rather than its members, is likely to be critical to the community's success. Iriberri and Leroy [2009] showed that owner-generated content, when fresh and interesting, plays a key role in the success of communities during their growth stage. Preece et al. [2004] investigated lurking behavior and found that moderation is critical to the overall success of online communities. Matthews et al. [2013] stated that although community leaders are critical for fostering successful communities, existing technologies rarely provide them with direct support.

2.2. Group and Community Recommendation

Community or group recommendations have been studied in many domains, such as TV programs [Masthoff 2004], travel [Jameson 2004], and cooking [Berkovsky and Freyne 2010]. Many studies focus on the challenge of representing the group's preferences based on the preferences of the individual members. Techniques range from aggregating the individual profiles into a single profile that represents the group and generating recommendations for the aggregated profile [Masthoff 2004] to generating recommendations for the individual members and aggregating the recommendations themselves [O'Connor et al. 2001]. Our algorithm uses the first technique by studying the aggregation of interest profiles of three types of community members: owners, active members, and regular members. Unlike typical group recommendations, which are provided to all members, our recommendations are targeted to specific owners, who serve as "proxies" to the community and decide whether to share with all other members.

In our previous work, we introduced in detail a set of algorithms for recommending content to community owners [Ronen et al. 2014]. The algorithms were based on the community's content (title and metadata), members (different subsets), and combinations of both content and members. That work focused on comparing *interest* ratings, given by the owners in a one-time experiment, both to them personally and to the community as a whole, using recommender systems criteria. It found that the group of active members was the most effective for increasing interest ratings and that it became even more effective when hybridized with content. In this work, we apply the same set of algorithms but focus on their effect on the community. The comparison among the algorithms is based on owners' sharing behavior, reflected both through a hypothetical question and through real action, and the impact of the recommendations on community activity over time. To simulate a more realistic scenario wherein recommendations are continuously shared, we use four recommendation rounds. As part of our analysis, we discuss the common and difference between the sharing behavior results and the rating analysis presented in Ronen et al. [2014]. Most prominently, we found that when it comes to sharing behavior, the content profile outperforms all member-based profiles (MBPs). We also examine the effect of new factors, such as community age and the presentation of related people and tags, on owners' sharing behavior. Finally, and perhaps most importantly, we show that the recommendations have a significant effect on community activity in the weeks that follow the recommendation rounds, both in terms of activity volume and in terms of active members.

2.3. Engagement in Social Media

In this work, we use recommendations to encourage content sharing by community owners. Several previous studies have used techniques to promote content contribution and engagement in social media. Freyne et al. [2009] showed that new users of an enterprise social network site, who were presented with recommendations of relevant people and content, contributed significantly more content than new users who did not get such recommendations. Cosley et al. [2007] used intelligent task routing of Wikipedia articles to editors and showed that it could substantially increase their contribution. Dugan et al. [2010] recommended topics for blog authors to post about. They found that posts created through their recommender system led to more interactivity and drove more traffic. Hilts and Yu [2010] investigated two recommendation strategies to encourage debates in a website devoted to climate change discussions. They showed that the choice of strategy for recommendation can lead to different kinds of engagement, such as high participation or polarized discussions. In our case, we focus on the community owners as the target population.

Other ways of increasing the engagement of social media users have been explored in the literature. In recent years, the domain of gamification has emerged as a way to incentivize users and increase engagement through the inclusion of playful elements, such as leaderboards and badges; Cronk [2012], Guy et al. [2015], and Farzan et al. [2008] developed a point system in an enterprise social media site to reward contribution and observed a substantial increase in activity, over the short term. Kraut and Resnick [2012] suggested spiking motivation of members to contribute through the way contributions are being requested and rewarded. Ludford et al. [2004] showed that indication of similarity and uniqueness of movie ratings could spark contributions. Visualizations have also been harnessed to increase engagement [Perer et al. 2011]. For example, Sun and Vassileva [2006] developed a motivational visualization encouraging social comparison; this was shown to yield a significant growth in participation and contribution to peer-to-peer online communities.

3. RECOMMENDER SYSTEM

3.1. Research Platform

Our research was performed over an IBM Connections (IC)[1] deployment within IBM. IC is a social media platform, which includes various social media applications, such as blogs, microblogs, wikis, discussion forums, activities, bookmarks, and shared files. Communities contain a subset of these applications that are tailored for use in the context of a community. A community can thus include a blog of its own, a wiki of its own, and so on. Communities can be public, invitation-only, or private. They define two types of users, *owners* and *members*. Owners have administrative privileges such as designing the look of the community, moderating and deleting content, and adding new owners or members. They do not have an official definition of a management role within the community. Consequently, there are owners who are not actively involved in the community, whereas regular members might take the actual leadership [Ehrlich et al. 2014; Matthews et al. 2014a].

[1]http://www-03.ibm.com/software/products/en/conn.

We measure the activity level of a community by counting the actions that were performed in the context of the community. These actions can be posting or commenting on a microblog in the community; creating, commenting on, or liking a blog post of the community; creating, commenting on, or liking a forum entry of the community; creating, editing, or liking a wiki page; sharing, downloading, or liking a file; adding a bookmark; or creating an activity in the community [Muller et al. 2012]. Members are not automatically updated about new content in the community. They need to actively follow a community by clicking the follow button on the community's homepage, which triggers updates either through an email digest or on their IC homepage. At the time of this study, the IC deployment included nearly 200,000 communities, of which about 120,000 were public or invitation only. We examined only public and invitation-only communities, since private communities were not accessible.

3.2. Community Interest Profile

Recommender systems have traditionally used one of two approaches: The content-based (CB) approach [Pazzani and Billsus 2007] generates recommendation based on items whose content and metadata (e.g., title, actors, or tags for movies) are similar to those already liked by the user. Collaborative filtering [Goldberg et al. 1992] recommends items users with similar preferences or tastes have liked, allowing more diverse recommendations. Hybrid methods [Burke 2002] combine the two approaches and often show improvement over each separately. In this work, we analogously experiment with a content-based approach, a member-based approach, and a hybridization of both. One of our key research questions is which of these approaches would generate recommendations most likely to be shared by owners with their communities. More specifically, is it better to recommend based on the community's content or based on the interest of its members? If the latter, then which subset of the members should be used? And finally, is it beneficial to hybridize content-based and member-based approaches?

With these questions in mind, we generated recommendations for a community based on an interest profile. We examined seven interest profiles that were, as mentioned, either member based, content based, or hybrid. MBPs represented the interest profile of the community members or a subset of these. The content-based profile (CBP) modeled the community's interests according to its title and metadata. The hybrid approach combined the interests of the members with the content of the community. For generating the profiles and recommendations, we used the algorithm described by Ronen et al. [2014]. In the remainder of this section, we describe this algorithm and the intuition behind it.

3.1.1. Member-Based Profiles. We examined three types of MBPs, each based on a different subset of the community members. The first two groups represent the two formal roles defined by the community (members and owners) and the third group represent active members of the community, whose importance was recognized in a recent study of enterprise online communities [Ehrlich et al. 2014]. Specifically, the *Members* profile was based on 50 random members (including owners) of the community or on all members of the community if there were 50 members or less. The random sampling was applied in order to avoid the high computational costs when involving many members.[2] The *Owners* profile considered all the owners of the community. The *Actives* profile was based on the set of members (including owners) who were active at least once within the community along its lifetime, according to the definition of a community's activity stated in the previous subsection.

[2]We also experimented with a profile of 100 random members and found it produced very similar recommendations; we therefore opted to use the less computationally intensive version of 50 members.

Our method for generating the MBPs builds on the method used for recommending social media items to individual users [Guy et al. 2010]. Individual member profiles consisted of *profile elements* that included related people, denoted by P, and related tags, denoted by T. In our experiments, we set $|P| = |T| = 30$, as in Guy et al. [2010]. For the community profile, we considered the individual profile of each of the members and aggregated them into a single community profile, which itself included 30 people (P_c) and 30 tags (T_c). The aggregation of multiple individual profiles into one community profile was based on the number of members whose profiles contained each profile element and the relative position (rank) of the profile element in each of these profiles. For tags, we also considered stemming and inverse document frequency. Below, we describe the method in more detail.

Let M denote the set of members an aggregated profile was based on. For each member $m \in M$, we computed an individual profile, denoted as $prof(m)$, which included the top $|P|$ related people and top $|T|$ related tags, ranked by their relationship strength to that member. Related people were calculated and ranked based on familiarity relationships reflected in social media, such as explicit "friending," wiki page co-editing, file sharing, and others, as well as similarity relationships, such as bookmarking of the same pages, usage of the same tags, membership in the same communities, and so on. Related tags included tags used by the member to annotate different entities as well as tags that were assigned to them by others within an enterprise people tagging application. Full details of the individual profile calculation can be found in Guy et al. [2010].

3.1.1.1. Related People. Given these individual profiles, the list of $|P_c|$ people to be included in the aggregated profile for M was determined according to the following scoring formula:

$$score(p, m) = \begin{cases} 2\,|P| - rank_{prof(m)}(p) & p \in prof(m) \\ 0 & p \notin prof(m) \end{cases}$$

$$score(p, M) = \begin{cases} \sum_{m \in M} score(p, m) & count(p, M) \geq 2 \\ 0 & count(p, M) < 2, \end{cases}$$

where $count(p,M)$ denotes the number of members in M that included person p in their individual profile, and $rank_{prof(m)}(p)$ denotes the rank of person p out of all $|P|$ people included in the individual profile of a member $m \in M$. The rank of the top person in the profile would be 0, the second person would get a rank of 1, and so on. Note that the measure $2|P| - rank_{prof(m)}(p)$ assigns the top person with a score that is almost double the score of the bottom person in the profile: $2|P| - 0 = 60$ versus $2|P| - (|P| - 1) = |P| + 1 = 31$, respectively, in our case. This is done to limit the influence of the rank within an individual profile up to a factor of 2. For example, a person who appears at the bottom of two members' profiles would get a higher score than a person who appears at the top of only one member's profile, with 62 versus 60. Ultimately, we summed these scores across all members in M, considering people who appeared in at least two member profiles, to make sure they had at least two different member "votes," and selected the top $|P_c|$ according to their score. Thus, the more members in M a person is related to and the stronger the relationship to them, the higher the chances of that person to be included in the aggregated community profile. Finally, we normalized all scores by the score of the top person in the profile.

3.1.1.2. Related Tags. The list of $|T_c|$ tags in the aggregated profile was calculated in a similar manner, with two adaptions addressing the need for stemming and for penalizing popular tags, which tend to be very broad and less meaningful. We first applied stemming [Manning et al. 2008] in order to merge similar forms of tags, such as "travel," "traveler," and "traveling". The score of a stemmed tag t for the aggregated profile of M was calculated according to the following formula:

$$score(t, m) = \begin{cases} 2\,|T| - min_{t' \in stem(t,m)}(rank_{prof(m)}(t')) & t \in prof(m) \\ 0 & t \notin prof(m) \end{cases}$$

$$score(t, M) = \begin{cases} idf(t) \cdot \sum_{m \in M} score(t, m) & count(t, M) \geq 2 \\ 0 & count(t, M) < 2, \end{cases}$$

where $stem(t,m)$ denotes the set of tags in the profile of a member m that convert into tag t after stemming. Analogously to the people case, $count(t,M)$ denotes the number of members whose profiles include t and $rank_{prof(m)}(t')$ denotes the rank of a non-stemmed tag t' out of the $|T|$ tags included in the profile of a member m. Finally, the inverse document frequency of a stemmed tag t, $idf(t) = ln(N/Nt)$, is computed as the logarithm of the ratio between the total number of documents in the system (N) and the number of documents tagged with at least one tag that converts into t after stemming (Nt). Similar to the vector-space *idf* score for terms [Manning et al. 2008], the *idf* score for tags penalizes popular tags, which are related to many documents. The total score of the stemmed tag was calculated by summing the scores over all members, for tags that appeared in the profiles of at least two members. The top $|T_c|$ tags with highest scores were then selected for the aggregated profile, with their scores normalized by the highest value. Intuitively, a tag would have higher chances of being included in the aggregated profile if it is related to more members in M, if the relationship to each of these members is stronger, and if the tag is generally less common.

3.1.2. Content-Based Profiles. The CBP considered the community's title (must be present for any community), description (83.7% of the communities had it), and tags (79.5% of the communities had tags),[3] all typically added by owners at community creation time. We used the $KL + TB$ measure [Carmel et al. 2012] to identify the most significant terms (of up to three words) in the extracted content. This method was previously found effective for term extraction from concise social media content [Carmel et al. 2012]. The method uses the Kullback–Leibler (KL) measure, which is a non-symmetric distance measure between two given distributions. In our case, we sought out terms, in their stemmed form, which maximize the KL divergence between the language model of the community's content and the language model of the entire community collection's content. On top of the KL statistical score, we applied a tag-boost (TB), which promotes keywords that are likely to appear as tags, based on a given well-tagged folksonomy. For this purpose, we used the folksonomy generated by the IC bookmarking application [Millen et al. 2006].

Ultimately, a community's content profile included all terms that had a $KL + TB$ score that was at least 30% of the maximum $KL{+}TB$ score of a term in that community. We experimented with various other thresholds, but found 30% to yield the best tradeoff between the overall number of extracted terms and their quality.

3.1.3. Hybrid Profiles. We hybridized each of the three MBPs with the CBP by considering both the people and tags included in the MBP and the terms included in the CBP. Accordingly, the *MembersContent*, *OwnersContent*, and *ActivesContent* profiles were

[3]We did not use the full content of a community's items, since our initial experimentation indicated they were often noisy or focusing on very specific aspects of the community.

defined, consisting of people and tags from the MBP and content terms from the CBP. We further describe how recommendations were generated for the hybrid profiles in the next section.

3.3. Recommendation Generation

Given a community profile, we generated recommendations by issuing a query containing the profile elements to a social search system [Ronen et al. 2009], similarly to the way it was done for an individual profile in Guy et al. [2010]. The social search system, which is built on top of Lucene [McCandless et al. 2010], indexes social media documents of different types, including blog entries, wiki pages, shared files, forum threads, activities, and bookmarks (see [Muller et el. 2012; Perer et al. 2011] for more details on each of these types). The system maps the relationships among these documents, related terms and tags, and related people, in a way that makes all types of entities both searchable and retrievable [Ronen et al. 2009]. For the task of producing recommendations, the query to the social search system included a combination of people, tags, and terms, whereas the results were documents that matched the query, ordered by their relevance score. Below we describe in more detail the queries and the calculation of the relevance score.

For the non-hybrid profiles, we retrieved the top 100 documents by issuing an OR query to the social search system. This query included all the profile elements as its arguments, each boosted with its corresponding score, calculated as explained in the previous section. For a profile that included people $p_1 \ldots p_u$ with scores $s(p_1) \ldots s(p_{u,})$ and tags $t_1 \ldots t_v$ with scores $s(t_1) \ldots s(t_v)$, we issued the following query:

$$q = (p_1 {}^\wedge s(p_1) \vee \ldots \vee p_u {}^\wedge s(p_u)) \vee (t_1 {}^\wedge s(t_1) \vee \ldots \vee t_v {}^\wedge s(t_v)).$$

The symbol "^" denotes the boosting factor.

For a hybrid profile, consisting of an MBP with people $p_1 \ldots p_u$ scored by $s(p_1) \ldots s(p_u)$ and tags $t_1 \ldots t_v$ scored by $s(t_1) \ldots s(t_v)$, and of a CBP with content-terms $c_1 \ldots c_r$ scored by $s(c_1) \ldots s(c_r)$, recommendations were created by issuing the following query to the social search system:

$$q = ((p_1 {}^\wedge s(p_1) \vee \ldots \vee p_u {}^\wedge s(p_u)) \vee (t_1 {}^\wedge s(t_1) \vee \ldots \vee t_v {}^\wedge s(t_v))) \wedge (c_1 {}^\wedge s(c_1) \vee \ldots \vee c_r {}^\wedge s(c_r)).$$

The query retrieved the top 100 documents that were relevant to at least one person or tag from the MBP and one content term from the CBP. In this way, we made sure that the returned documents matched both parts of the hybrid profile.

Upon receiving a query q, the relevance score of a document d in the social search system was calculated as follows:

$$RS(d,q) = e^{-\alpha\tau(d)} \cdot \left[\beta \sum_{i=1}^{u} s_q(p_i) \cdot w(d, p_i) + \gamma \sum_{j=1}^{v} s_q(t_j) \cdot w(d, t_j) + (1 - \beta - \gamma) \sum_{k=1}^{r} s_q(c_k) \cdot w(d, c_k) \right].$$

Notice that the third and final element of the summation is only relevant for hybrid profiles, otherwise it is disregarded. In the equation, $\tau(d)$ denotes the time in days since the creation date of d; α is the time-decay factor, used to promote fresher documents (set in our experiments to 0.025, as in Guy et al. [2010]); β and γ are the parameters that control the relative weight among people, tags, and content terms. In our experiments, we set both to 1/3, giving equal importance to all ingredients; $S_q(p_i)$, $S_q(t_j)$, and $S_q(c_k)$ are the scores of the respective profile elements, given as part of the query q; and $w(d,p_i)$, $w(d,t_j)$, and $w(d,c_k)$ denote the relevance score of the document to the person, tag, or content-term, as calculated by the social search system (see more details in Guy et al. [2010] and Ronen et al. [2009]).

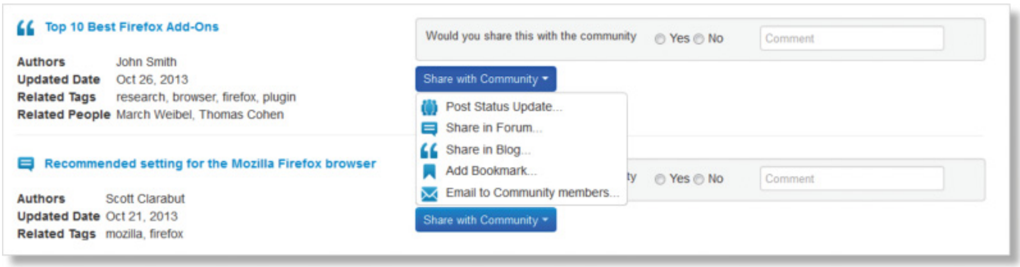

Fig. 1. Two sample recommended items in the survey.

Ultimately, we selected the top 10 items for recommendation in our survey after applying the following two steps over the 100 retrieved documents: (1) *filtering*: documents that were already published in the community were filtered out; (2) *diversifying*: in order to promote diversity across document types (blog entry, wiki page, and so on), we used the type as the first sorting criterion and the relevance score only as a secondary criterion. Therefore, we first took the top document of each type, if such existed among the top 100 documents, and ordered these by their relevance score. We then took the second of each type, if one existed, and ordered this group by the relevance score, and so forth until we reached 10 items (documents) in total. Finally, we randomized the order of all 10 recommendations.

4. EXPERIMENTAL SETUP

4.1. Owner Survey

Our evaluation was based on a survey of community owners that included four rounds of recommendations. Each round was only sent to the owners who had responded to the previous round. Rounds were 2 weeks apart over a period of 6 weeks in total. The owner received a personal email invitation that included a link to the community and a link to the online survey of the particular round. The email described the survey and thanked the owner for participating in previous rounds, when relevant (full wording is provided in Appendix). We opted for a multiple-round survey, since we wanted to simulate, even if roughly, a real-life situation wherein recommendations are available on a regular basis, and inspect the effect on communities over time.

In each round of the survey, owners were presented with a set of recommended items. The first round included 11 items: 10 of them were generated based on one of the seven profiles described in the previous section, whereas an extra item was randomly drawn from the social search index to serve as a weak baseline. The position of the random item within the list of 11 items was randomly selected for each community. Rounds 2 to 4 included five recommendations, all based on the same profile as in the first round. We opted to present fewer recommendations in the later rounds in order to encourage participants to take part in additional rounds by reducing the load.

Figure 1 illustrates how recommended items were presented in our survey. Each recommendation included the title of the item as a link to its IC page. An icon indicated the type of the item, i.e., forum entry, blog post, wiki page, activity, file, or bookmark. Additionally, the author's name, last update date, and the item's related tags and related people were presented. Related people included individuals other than the author, who performed some action on the item, such as commenting, sharing, or liking [Guy et al. 2016]. Related tags included tags that the item had been annotated with, if any existed. For each item, the owner was requested to indicate whether s/he would share the item with the community (yes/no answer). This was a hypothetical question

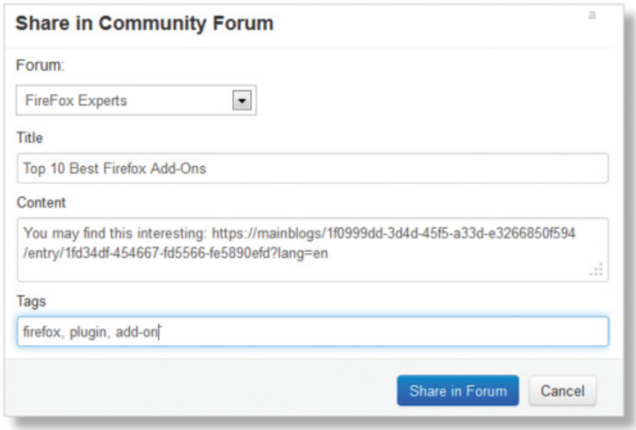

Fig. 2. Sharing popup for a forum entry.

that did not trigger any real action. Additionally, optional free-text comments could be added per a specific recommendation and at the end of the survey.

The owner could also actively perform a sharing action by clicking the “Share with Community” button, which enabled sharing with the community by creating a microblog message, a forum entry, a blog post, or a bookmark, which was then automatically published in the community. The owner could also choose to share by sending an email to all community members. Upon selecting one of the five sharing actions, a popup window would appear with a default text as the content, including the title of the recommended item and a link to it. The owner could edit the text before it was published in the community (or sent via email). Figure 2 shows a sample popup for a forum entry. Performing a sharing action was optional and, since this was a survey setup, we did not expect a huge amount of actions to be executed. Nevertheless, we hoped to conduct some experimentation with real actions that have an immediate influence on the community.

The first round of the survey was sent to owners of communities that had at least five members and at least two owners, to make sure they have a sufficient basis for producing the MBPs. These accounted for 83% of all public communities. In addition, we only considered communities for which six activities were performed during the two months preceding the survey, to make sure we focus on communities that are still relevant to their members. This left us with a sample of nearly 3,000 communities. Each community in this sample was randomly assigned to either one of the seven profiles described before or to an eighth group that served as the control group and did not receive any recommendations. We sent an invitation to participate in the survey to at most three owners of each community. We preferred active owners, if existed, since we reckoned that they would be more eager to take part in the survey. We also made sure an owner would get at most three surveys for different communities. As stated before, for subsequent rounds, we only invited owners who responded to the previous round for the respective community.

4.2. Research Tasks and Hypotheses

Our experiments were designed to address three key research tasks, each with its own hypotheses, as follows.

RT1: Compare the effectiveness of member-based, content-based, and hybrid profiles for recommendation sharing by owners

Table I. Participation Across the Four Survey Rounds

Round	1	2	3	4
Participants	1,100	497 (45.2%)	264 (53.2%)	184 (69.7%)
Owners	1,053	488	258	183
Communities	903	447	242	177

H1.1 Similarly to the case of ratings [Ronen et al. 2014], hybrid profiles outperform content-based and MBPs.

H1.2 All community profiles are significantly more effective than the weak random baseline.

RT2: Examine the effect of different characteristics of the community and the recommended items on recommendation effectiveness

H2.1 Community's age and size have an effect on recommendation effectiveness.

H2.2 Recommended item's type and presentation of related people and tags have an effect on recommendation effectiveness.

RT3: Examine the effect of recommendations on community activity in the weeks that follow, relative to the control group who received no recommendations

H3.1 The number of activities within the community grows as a result of the recommendations.

H3.2 The number of distinct contributing members grows as a result of the recommendations.

In addition to these three research tasks, we set out to explore, using both quantitative and qualitative methods, further aspects of recommendation sharing by community owners, in order to gain a broader understating of this process. These aspects include owners' motivations for sharing content, how they behave over time, what tools they use to share content, and which communities this type of recommendation is more suitable for.

Our result analysis follows the three key research tasks. First, we analyze the sharing behavior across the seven profiles as reflected both through the hypothetical survey question (would you share?) and real actions taken through the survey. We then analyze the effects of community characteristics (age, size) and item characteristics (type, related people and tags) on sharing behavior. Finally, we compare community activity in the eight weeks following our survey, for communities whose owners received recommendations versus the control group. All the results are described in detail in the following section.

5. RESULTS

Table I summarizes the number of responses we received in each round and the number of owners and communities covered by these responses. We refer to each response as a "participant" in the survey, even though participating owners could have multiple responses that corresponded with multiple communities. In parentheses is the percentage of participants from the previous round who also participated in the current round. It can be seen that there is a substantial decrease in participation from one round to another, which gets milder for later rounds: Although less than 50% of round-1 participants continued to round 2, almost 70% of round-3 participants continued to round 4. Overall, only 16.72% of round-1 participants continued to round 4.

In the survey's comments, many participants expressed appreciation to the idea of sharing content recommendation with their community and pointed out various motivations. For example, one owner wrote: "*This should help educate team members, especially new ones, about relevant projects, resources, and ideas*" and another commented: "*sharing these links helped me show the community I care [. . .] over time I hope*

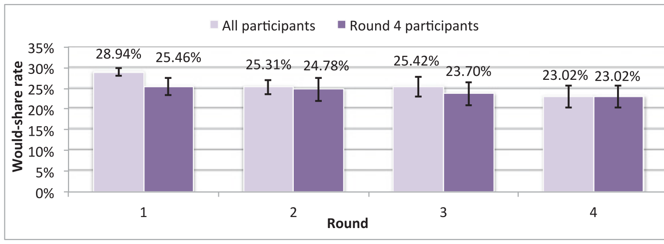

Fig. 3. Average would-share rates across the four rounds.

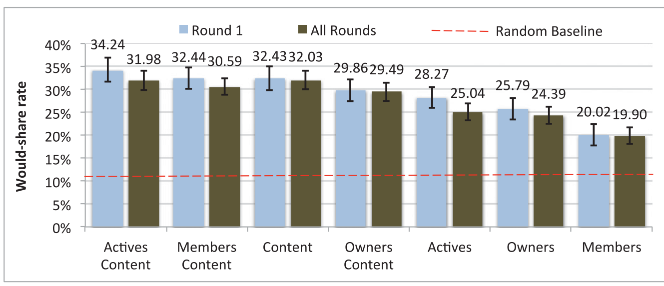

Fig. 4. Average would-share rates across the seven profiles.

it can create more sense of community". Several owners mentioned that they would like to see the recommendations featured on the community's page on a regular basis, e.g., "*at least on a monthly basis.*" Another owner summarized her community needs: "*the community wants to know (a) general trends and issues related to <topic> globally; (b) products and services available in the market; (c) what IBM has to offer; (d) client opportunities and (e) areas of IBM considering <topic> opportunities"*.

5.1. Would-Share Rates

The overall *would-share rates*, i.e., portion of recommended items for which owners indicated they would share them with the community, were 27.07% across all rounds and all profiles. Figure 3 shows the rates across the four rounds, both for all survey participants and for the 184 participants who took part in all four rounds. The error bars represent 95% confidence intervals for the would-share rates.[4] Overall, there is a slight decrease from one round to another: Possibly because owners become more selective when the recommendation process becomes routine. For all participants, there is a noticeable drop between the rates in round 1 and round 2. For round-4 participants, there is a small consistent decrease across all rounds. The fact that the decrease is mild shows that despite the growing challenge, recommendations remain effective also by the fourth round. It is also noticeable that would-share rates by round-4 participants in the first three rounds, and especially in round 1, are lower than the average rates for all participants, indicating that those who continued to the last round were not necessarily the ones who found the recommendations most suitable to share with the community in the first place.

Figure 4 shows the average would-share rates across each of the seven profiles, for round-1 participants and for aggregated participants of all rounds. The dashed line marks the random baseline, i.e., the average would-share rate of random items in

[4]In all figures, error bars represent 95% confidence intervals based on the adjusted Wald method [Brown et al. 2001].

Fig. 5. Sharing action distribution by type.

round 1, which was found to be 10.87%. In the first round, the results are partially similar to interest rating results reported in Ronen et al. [2014]: Among MBPs, *Actives* has the highest rates and *Members* has the lowest; hybrid profiles yield higher rates than their pure member-based counterparts; and the *ActivesContent* profile achieves the highest rates of all. The most noticeable difference from interest rating results, however, is that the pure *Content* profile yields high would-share rates, which are second only to the *ActivesContent* profile (and equal to *MembersContent*). It appears that when it comes to possible sharing with the community, the specific content of the community plays a more central role and owners are more willing to share items that directly relate to it.

Welch ANOVA for the first round indicates that would-share rates across the seven profiles were significantly different, $F(68{,}562.62) = 14.02$, $p<0.001$. Games–Howell post-hoc comparisons indicate that the average rates of the *Members* profile were significantly lower than for all other profiles; rates of the *ActivesContent*, *MembersContent*, and *Content* profiles were also significantly higher than for the *Owners* profile; and *ActivesContent* was the only profile with significantly higher rates than for all pure MBPs.

Inspecting the aggregated results across all four rounds, a general decrease in rates can be observed compared to round 1. The *Content* profile, however, suffers a very minor decrease and becomes equal to the *ActivesContent* profile. It appears that over time, owners tended more strongly towards sharing items that related to the community's content and thus the *Content* profile performed best for later rounds (3 and 4).

Overall, our hypothesis *H1.1* is rejected: The would-share results are not identical to the rating results reported in Ronen et al. [2014], where profiles that contained signals from active members were the strongest. In contrast, we have seen here that profiles that contain content signals perform better for sharing recommendations. Our hypothesis *H1.2* is supported: All profiles we experimented with perform better than the weak random baseline.

5.2. Sharing Actions

Overall, 1,033 sharing actions were carried out in our survey across all four rounds, over 7.23% of all recommended items. These actions were performed for a total of 340 communities (37.65% out of all) by a total of 354 owners (33.62%). For 213 communities (23.59%), two or more actions were performed and for 64 communities (7.09%), 5 or more actions were performed. The maximum number of actions per community was 21. Figure 5 shows the action distribution by type. Sharing the recommended item as a bookmark was clearly the most popular action with almost 50% of all performed actions. We note that, in general, adding a bookmark is not a very popular action in a community and accounts for less than 10% of all add-create-share actions in a community. Apparently, bookmarking was perceived as the most suitable manner to share a new content piece, possibly as it does not involve any additional content from the sharer. One of our participants noted: "*bookmarking was the fastest and most natural way to share the content.*" Sharing through a forum entry was the second most popular action. Indeed, forum activities are the most popular among add-create-share actions in communities and account for nearly 25% of them. One owner explained: "*I shared as a forum entry so I can add some level of description of what the link provides for the*

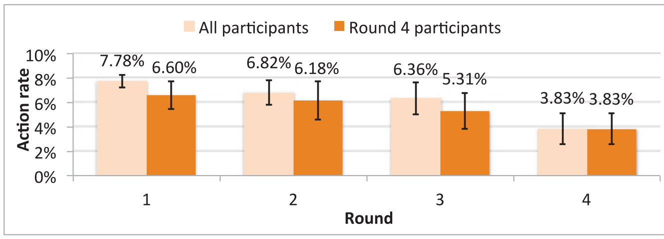

Fig. 6. Sharing action rates across the four rounds.

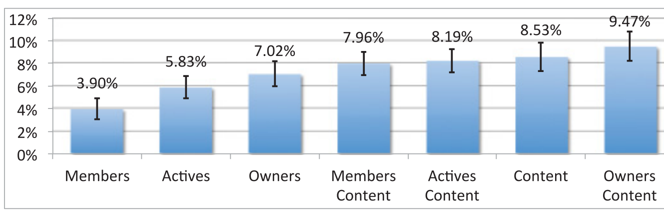

Fig. 7. Action rates across the seven profiles.

community". Sharing the item through a blog or microblog was less popular, possibly as they are typically characterized by more personal text. Email was very rarely used, likely as it was more intrusively "pushing" the item to the recipients.

In several cases, owners explained that the specific nature of the community was not suitable for sharing recommended content. For example, one commented: "*My community is very pinpointed [...] contains specific actions—files needed by the managers. [...] Other content would add confusion, when my community's purpose is to provide the material they need + instructions*". Another wrote: "*This community was created for team members to report their status [...] Dialog and news are shared in another community.*" A few participants mentioned they wanted better control on how the recommendations would appear in the community, as one noted: "*I see the options for sharing content, but I am not sure where these should show up or exactly how they are presented. Probably need to understand more about how to ensure they go to a place that is relevant, but does not conflict with the more organized approach we are trying to take.*"

Figure 6 shows the *action rates*, i.e., the portion of recommended items for which a sharing action was performed, across the four rounds, both for participants of all rounds and for round-4 participants. Similarly to the would-share rates, action rates dropped from one round to another. The most noticeable decrease was from round 3 to 4. Although would-share rates only slightly dropped from round 3 to 4 for round-4 participants (Figure 3), sharing actions had a sharper decrease, indicating that the "fatigue effect" is stronger when it comes to real actions, compared to hypothetical sharing. Additionally, it can be seen that those who continued to round 4 performed less sharing actions than the general average per round, again implying that the continuation to later rounds was not necessarily dependent on the perceived quality of recommendations.

Figure 7 displays the sharing action rates by profile type. Results should be considered with care due to the relatively a low amount of overall actions. Generally, we can see the same trend as with the would-share rates: Hybrid profiles and *Content* have higher rates, whereas pure profiles and especially the *Members* profile have lower rates. *OwnersContent* is the top profile when it comes to action rates, followed by *Content,* and then the two other hybrid profiles. It seems that when it comes to real rather than

Table II. Characteristics of Participating Communities

	Size		Age	
	Avg	Median	Avg	Median
All	316.47	63	574.04	418.5
Round 4	393.71	69	582.1	442

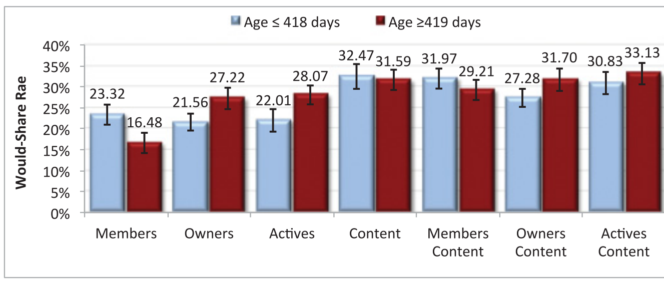

Fig. 8. Would-share rates for old versus young communities.

hypothetical sharing, the owner-based profiles (both pure and hybrid) are better than the active-based ones. In all, these results support our rejection of hypothesis *H1.1*: The CBP is superior to all MBPs and is comparable to the hybrid profiles.

5.3. Community Characteristics

To examine hypothesis *H2.1*, in this subsection, we inspect the would-share rating results with regard to two key characteristics of the participating communities: size (number of members) and age (days between creation date and first round date). Table II displays the average and median size and age for communities whose owners participated in the survey across all rounds and also for those who took part in round 4 (i.e., participated in all four rounds). The characteristics of the latter are not very different from those of all communities, although they are slightly larger and older. For all participating communities, there was a small positive correlation between size and age ($r = 0.22$, $p < 0.001$).

As Table II indicates, the median age of participating communities was 418.5 days (about 14 months). Figure 8 shows the would-share rates across the seven profiles for communities that were created 419 days before the survey or earlier (old communities) versus communities that were created within the 418 days preceding the survey (young communities). Generally, old communities received significantly higher rates than young ones: 28.89% versus 26.67%, (one-tailed unpaired t-test, $p < 0.05$). It appears that owners of older communities, which are typically more established, are more willing to share content with their members. Inspecting the results by profile type reveals that the profiles that are based on regular members (*Members* and *MembersContent*) received higher rates for young communities, whereas the profiles that are based on a subset of the members (*Owners*, *Actives*, *OwnersContent*, and *ActivesContent*) received higher rates for old communities. Apparently, as a community becomes older, its regular member list becomes noisier, whereas the groups of owners and active members are better established and provide a more effective basis for recommendation, particularly when combined with content.

Figure 9 shows the would-share rates for communities with 63 members or less (small communities) versus communities with over 63 members (large communities), across the seven profiles. Generally, would-share rates were significantly higher for large communities than for small communities: 28.42% versus 26.98% (one-tailed unpaired

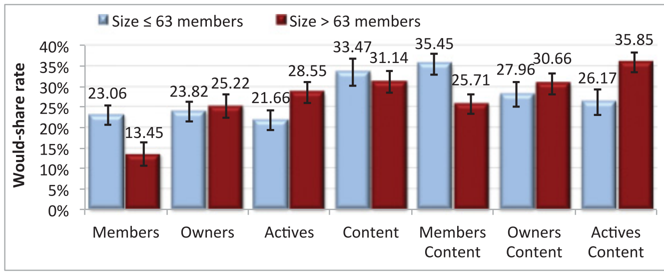

Fig. 9. Would-share rates for small versus large communities.

t-test, $p < 0.05$). This may be due to the fact that smaller communities often represent narrower, more specific, and even ad-hoc topics and thus fewer recommendations will fit their profile. Inspecting the results by profile type indicates that the *Members* and *MembersContent* profiles receive higher rates for small communities, whereas *Owners* and *OwnersContent*, and especially *Actives* and *ActivesContent,* received higher rates for large communities. It appears that for larger communities, recommending based on a random subset of the members is less productive, whereas the smaller group of active members becomes more representative and effective for recommendation, especially when combined with content.

Overall, we found support for our hypothesis *H2.1*—both the age and size of a community are correlated with the willingness of owners to share recommendations. We note, however, that this does not necessarily mean that age and size influence willingness to share, as the other direction is also possible: It could be that having an owner who is more willing to share information leads to the community surviving longer and growing faster.

In the survey's comments, several owners mentioned the recommendation topics were too broad for the focused topic of their community: "*This is a very specific community and we're trying to keep a high 'signal to noise ratio'. Some of these recommendations are interesting but are not directly related to* the goals of the *community—improving collaboration between <dept1> and <dept2>*." On the other hand, some owners mentioned their community is very broad and therefore narrow topics might be irrelevant for many members. For example, one commented: "*The community that I maintain has 10K members and it is problematic to share links that are narrow in scope [...] for example, because of regional differences it is tricky to propose information that is geo-specific.*"

5.4. Item Characteristics

Overall, we found support for our hypothesis *H2.1*—both the age and size of a community are correlated with the willingness of owners to share recommendations. We note, however, that this does not necessarily mean that age and size influence willingness to share, as the other direction is also possible: It could be that having an owner who is more willing to share information leads to the community surviving longer and growing faster.

In this subsection, we focus on hypothesis *H2.2*. Figure 10 shows the distribution of would-share rates and action rates for the different types of recommended items, taking into account all survey rounds. In parentheses is the portion of items of each type out of the overall set of recommended items. It can be seen that both rates are rather similar across the different types, with a noticeable difference in favor of bookmarks. Overall, these results indicate that mixing item types for sharing with the community is appropriate, as all types produced rather similar results. Bookmarks and blogs, which

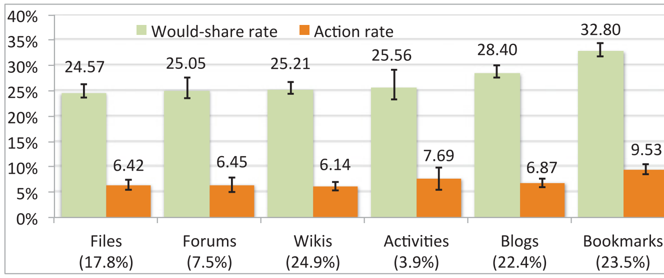

Fig. 10. Would-share and action rates by item type.

Table III. Would-Share and Action Rates Based on the Display of Related People and Related Tags

	None	People	Tags	Both
Portion	49.83%	16.28%	42.81%	8.92%
Would-share rate (95% CI)	23.81% (±1.01)	30.48% (±1.91)	30.48% (±1.18)	31.42% (±2.61)
Action rate (95% CI)	6.36% (±0.58)	7.77% (±1.12)	8.11% (±0.7)	7.58% (±1.5)

have the higher rates, were also among the more commonly recommended types. Wikis were most commonly recommended but did not yield better results, thus a negative boost may be considered for them when scoring recommendations.

Table III shows the would-share and action rates for items with no related people and no related tags displayed, only one of these, or both. It can be seen that items with related people, tags, or both, had substantially higher rates than items with no related people and tags. Guy et al. [2010] studied explanations in the form of people and tags that relate to both the item and the user. They found that only people-based explanations increased interest in recommendations for individuals, whereas tag-based explanations had no effect. In contrast, in this work, the recommender displayed people and tags related to the item but not necessarily to the user. We found that showing either of them increased the willingness of owners to share recommendations. It appears that with regard to items to be shared with their community, explanations are of high importance for owners. One of our participants wrote: "*More info—description, people, tags—helps me understand the value to the community [...] I wouldn't share something I don't understand, and I probably won't spend time to learn more except for rare cases where something in the title makes it especially promising.*"

Our hypothesis *H2.2* is supported by the results presented in this subsection: Recommendation effectiveness increases for certain types of items (bookmarks and blogs) and for items accompanied by the presentation of related people and related tags.

5.5. Community Activity Effects

In this key (and final) part of our evaluation, we examined the effect of our survey on community activity, as previously defined. Community activity is a widely used measure for community success [Iriberri and Leroy 2009; Matthews et al. 2014b; Preece and Maloney-Krichmar 2003]. We focused on the eight weeks preceding our survey (the survey started on August 8th) versus the eight weeks that followed. During the latter period, rounds of the survey continued every two weeks, until September 19th. In this analysis, we disregarded any activity that was performed by using the sharing action in our survey. As mentioned in Section 4, a random eighth of our sampled communities

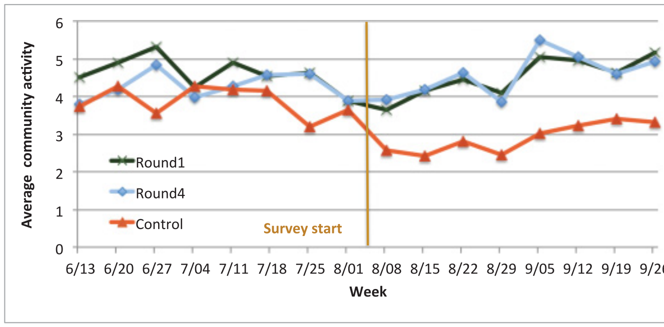

Fig. 11. Community activity before and after the survey.

were assigned to a control group who did not receive any treatment, i.e., none of its owners received recommendations. This group included 366 communities in total.

Figure 11 shows the average activity level for the communities that took part in round 1 (i.e., all survey communities for which at least one owner responded, see Table I) and round 4, compared to the control group. Note that round-1 communities and round-4 communities are not independent groups: The latter is a subset of the former. Each data point represents the number of actions within the community along the week that started with the date marked on the x-axis. For example, the 8/15 data points represent the number of actions along the week between August 15th and August 21st inclusive. It is important to note that we compare the behavior of communities that participated in the survey to those in the control group. There is no point comparing the volume of activity in the same group before and after the survey, since these could be affected by external factors. Specifically, the survey started during summer vacation time in many countries; thus activity was generally low until the end of August and started to rise back in September.

It can be seen that before the survey started, all three groups had a rather similar level of activity, whereas after the beginning of the survey, the activity level of both round-1 and round-4 communities became substantially higher than the control group. This difference was consistent and stable across all eight weeks that followed the survey start.

Table IV presents these results in an aggregated way (average and median) for the entire period of eight weeks before and after the survey. The t-test values indicate the results of a one-tailed unpaired t-test between the control group and round 1 or round 4, respectively, including the statistical significance (p), degrees of freedom (df) and Cohen's effect size (d). The average number of actions for the control group dropped from 31.05 (stdev: 34.14, median: 15) in the 8 weeks before the survey to 23.26 (stdev: 26.76, median: 11) in the 8 weeks that followed. As mentioned before, this is due to the "vacation effect" during August. In contrast, for round-1 communities, the overall number of actions remained rather steady: 36.96 (stdev: 35.22, median: 17) before versus 36.22 (stdev: 39.45, median: 16) after. For round-4 communities, the activity even increased in the eight weeks after the survey from 34.16 (stdev: 34.98, median: 17) to 36.69 (stdev: 40.81, median: 18), in spite of the vacation effect spotted for the control group. Overall, although before the survey there was no significant difference between the activity of round-1 and round-4 communities to the control group, after the survey, both were significantly higher than the control group.

The significant increase in activity observed for round-1 communities was only slightly less prominent than for round-4 communities. This may imply that one round of recommendations has a good enough outcome; possibly as it already creates the

Table IV. Average, Median Activity and Unique Users Within Communities Before and After the Survey

	8 weeks before			8 weeks after		
	Average	**Median**	**t-test**	**Average**	**Median**	**t-test**
Activity						
Control	31.05	15		23.26	11	
Round 1	36.96	17	p>0.05 (df=1,267, d=0.17)	36.22	17	p<0.01 (df=1,267, d=0.39)
Round 4	34.16	17	p>0.05 (df=541, d=0.09)	36.69	18	p<0.01 (df=541, d=0.4)
Round 1—No Action	36.31	17	p>0.05 (df=937, d=0.14)	35.13	16	p<0.01 (df=937, d=0.35)
Activity Excluding Invited Owners						
Control	18.52	11		16.03	9	
Round 1	23.4	13	p>0.05 (df=1,267, d=0.19)	24.72	13	p<0.01 (df=1267, d=0.35)
Round 4	22.33	13	p>0.05 (df=541, d=0.15)	24.62	14	p<0.05 (df=541, d=0.36)
Unique Users						
Control	7.09	4		6.09	3	
Round 1	8.82	5	p<0.05 (df=1267, d=0.12)	8.66	5	p<0.01 (df=1267, d=0.21)
Round 4	9.17	5	p>0.05 (df=541, d=0.14)	8.93	5	p<0.05 (df=541, d=0.23)

effect of owners' increased awareness and provides them with means for sharing content with their community. Yet, it should also be noted that the number of recommendations in round 1 was double the number in any other round and all owners of round-1 communities received at least two emails—one for round 1, which they accepted, and one for round 2, since they participated in round 1 (they would not have received the second email in a one-round setting). Further research should be conducted to examine the need for continuous recommendation rounds versus one periodic recommendation round (e.g., every three months).

One can assume that the increase in community activity stems from the sharing actions enabled in our survey: although sharing actions themselves were not counted as part of the activity, they may have prompted additional contributions. To further explore this, we examined the activity level of round-1 communities for which no action was performed through our survey (63.5% of round-1 communities). Results, shown in Table IV, indicate that for these communities, the level of activity after the survey was also significantly higher than for the control group, even though before the survey it was insignificantly higher. The decrease in activity level after the survey is slightly higher than for all round-1 communities, indicating that there was indeed a slightly stronger effect on activity for communities in which action was taken through our survey. Yet, this difference is small and evidently the effect also exists for communities whose owners participated in the survey but did not take any active action. This finding implies that the effect of our survey on community activity spans beyond the sharing action itself. It could be that owners also took action at a later time or that the recommendations increased their awareness of the community and led them to dedicate more attention to it.

We also set out to explore to what extent the increase in community activity is caused by the participating owner(s) versus by other members who were not directly exposed to the survey. To this end, we considered the activity performed only by members or owners who were not invited to participate in the survey. This was done to eliminate any direct effect of the survey on activity: For example, owners who received the invitation to the survey, even if they did not respond, might have been reminded of the community and triggered to contribute. The middle part of Table IV shows the results for this portion of the activity. It can be seen that the effect observed for the general activity remains similar: Activity decreased for the control group during the eight weeks following the survey, whereas it slightly increased for both round-1 and round-4 communities during these weeks. The differences between both round-1 and round-4 communities and the control group before the survey were not significant but became significant afterwards. Overall, we see that the increased community activity relative to the control group is caused not only due to the activity of owners who were exposed to the survey, but also by contributions of other members.

Table IV also presents the average number of unique users who performed at least one action within the community during the eight weeks before and after the survey started. Although there is a clear decrease in the number of users for the control group, there is only a very minor decrease for both round-1 and round-4 communities. These findings indicate that the effect of our survey on community activity is also reflected in the overall number of active individuals.

It can be seen that the activity and number of unique users for the control group were lower before the survey than for round-1 and round-4 communities (even if not significantly in most cases). We assume that this is due to the fact that the owners who opted to participate in the survey represent communities that are more active than the average. Indeed, when checking the activity level for round-1 communities whose owners did not participate, we found it was similar to the control group and even slightly lower.

Breaking down these results by profile type reveals that the effect was similar across all seven profiles. *OwnersContent* yielded the highest increase in activity, which may be explained by the fact that it is also the profile that triggered the highest portion of sharing actions.

Overall, hypotheses *H3.1* and *H3.2* are supported: Our recommendations increase both the number of activities and the number of unique active members in the eight weeks that follow our survey, in comparison to the control group. This increase was not proved to be a result of the sharing action in the survey; moreover, we found a similar effect for communities for which no sharing through the survey has occurred.

6. DISCUSSION, LIMITATIONS, AND FUTURE WORK

We reported a broad set of results about the willingness of owners to share content with their community through recommendations, their actual sharing behavior, and the effect of recommendations on community activity over time. We found that recommendations based on the community's content (title, summary, tags) are more effective than in previously reported results that examined perceived interest but not willingness to share [Ronen et al. 2014]. When it comes to the act of sharing an item with the community, its relation to the content becomes more central to owners' eagerness to share. The effectiveness of community's content becomes even more important in later recommendation rounds: Although in the first round the hybrid *ActivesContent* was the most effective profile, for later rounds, pure *Content* was the most effective.

Throughout our experiments, the *Members* profile consistently achieved the lowest results among all profiles, which indicates that basing recommendations purely on random members of the community is not desirable. When hybridized with the content of the community, however, the profile was shown to become much more effective. We note that our measurement is solely based on accuracy (i.e., percentage of items owners would share) and does not take into account other factors that may influence the overall value of recommendations, such as diversity and serendipity [McNee et al. 2006]; these may be supported by MBPs through the addition of people and tags that are not directly related to the community.

Further bisecting the results by community age and size revealed that the *Content* and *MembersContent* profiles were particularly effective for younger and smaller communities, whereas the *ActivesContent* profile was more effective for older and larger communities. It could be that as communities grow in size and age, the profiles that are based on pure content and/or regular members become noisier, and hybridization with a presumably more mature group of active members becomes more essential for providing recommendations that owners would share.

In our survey, we enabled owners to take real action and immediately share an item with the community. Overall, over 7% of the recommended items were shared with the communities, resulting in a total of over 1,000 sharing actions made through our survey. Segmenting these actions by profile type reinforces the superiority of profiles that utilize the content over the pure MBPs.

We found a substantial difference between the would-share rates and the real action rates. We conjecture that this gap is mostly due to the survey setup: The survey was not introduced as a tool that would also allow real action; the sharing action was an optional feature in the survey, which is not the familiar way of taking action in IC. It could be that owners wanted to share using the regular mechanisms in IC, with their own unique user experience and functionality. Additionally, there is likely to remain a gap between an indication of willingness to share in a survey and taking real action in practice; would-share rates may therefore be regarded as a form of an upper bound for taking real sharing action. This is in accordance with our expectation as stated in Section 4: We did not anticipate action rates to be similar to the “would share” rates but were rather hoping to experiment with a small amount of real actions, in addition to the main survey responses.

Bookmarks were the most successful type of recommended item, both with regard to willingness to share and to actual actions. Additionally, sharing items as bookmarks within the community was the most popular type of sharing action, despite the generally low popularity of bookmarks as a community content type. These results indicate a potentially important role social bookmarking [Millen et al. 2006] may play for recommending content to communities through owners, both since bookmarks are a good source for recommended items and since they serve as a good mechanism to share recommended items with the community.

Our survey included four rounds of recommendations with two weeks between each pair of rounds. Our number of participants declined from 1,100 in the first round to 184 in the fourth round, with over half of the participants dropping out before round 2. The participants who continued through all four rounds were not necessarily those who rated the recommendations higher or shared more items in the earlier rounds, indicating that accuracy is not the sole factor when it comes to keeping the owner engaged in sharing the recommendations. There is a mild decrease in willingness to share over rounds, possibly as owners get used to the process and raise their satisfaction bar. The decrease in sharing actions is more noticeable and points to a potentially bigger challenge in keeping recommendations effective over time. Incorporating relevance feedback and accounting for more diversity (e.g., across the item’s author or topic) can help maintain the effectiveness of recommendations over time. Displaying the impact of recommendations on communities over time can also help motivate owners to continue and share recommended items.

Our results show a clear impact of the survey on community activity in the eight weeks that followed the beginning of its first round, compared to the control group. This can be observed both in terms of the total number of contributions and the number of unique members who were active. These results are encouraging and show that the type of recommendation we propose can not only trigger potential (and real) sharing by the owner, but also contribute to the liveliness of the community over time. We found that the positive effect on community activity existed for all participating communities and not only those for which the owners took active action in the survey. This implies that the recommendations have a broader effect that may stem from the owners’ increased awareness of the community and of relevant content that may be shared with its members. Moreover, our analysis found that the effect on community activity spanned beyond the owners who were invited to participate and included other members who were not at all exposed to our survey.

6.1. Limitations

In our research for analyzing the sharing behavior, we did not use a complex model that incorporates all community and item characteristics, but rather inspected each of these factors separately. We did not apply repeated ANOVA, due to the substantial decrease in owners' participation from round to round. When reporting round-1 only results, we used Welch ANOVA with Games–Howell post-hoc comparisons, since there was no homogeneity of variances. Throughout the analysis, 95% confidence intervals are plotted based on adjusted Wald's test, which fits our data as each item is treated as an independent test with a possible outcome of 0 (no share) or 1 (share). We believe that the presented results support our hypotheses as phrased in this paper. Future research could examine more advanced models to gain a finer-grained understanding of the parameters to use for recommendation.

In this work, we considered only the community's title, summary text, and tags for its content profile. Content originating from the community's social media applications, such as blog posts, forum entries, wiki pages, or microblog messages, was not taken into account. Future work should examine whether these data could be used to further enhance the CBP. Our initial experimentation indicated that it is too noisy (in addition to being computationally intensive), but further methods, such as considering only the title and tags of these content items, or applying text summarization techniques, should be considered.

One of the core challenges for recommender systems is the cold start problem [Schein et al. 2002]. In truth, the methods proposed here are not applicable for brand new communities that do not yet have content or members. Other incentivizing techniques, such as gamification, can be used to spark initial contribution and then combined with recommendation methods, when the community already has some members and activities.

We note that the current mechanisms in IC to draw the attention of community members to new activity are "soft": Members need to actively choose to follow the community in order to get content notification updates. With more advanced mechanisms to update members about new content, engagement has the potential to further increase as a result of content sharing by the owner and the sharing actions' impact can be amplified [Guy et al. 2012].

Previous work examined the recommendation of content to create in order to boost participation in social media. For example, Cosley et al. [2007] introduced SuggestBot, which recommended Wikipedia volunteers with articles to edit. It should be noted that the target audience in this case was a group of volunteers with higher commitment to make contribution. In a closer environment to ours, Dugan et al. [2010] introduced BlogMuse, a system that suggested topics for blog authors to post. They found that these recommendations did not increase the total number of authored posts and speculated that posts created through BlogMuse replaced posts that would have been created otherwise. In our case, due to the generally low engagement level of community members, we opted to use the owners as the target population for recommendation, in the hope that their sharing of existing content with the members will increase contribution. Our results show that this method indeed increases participation in the following weeks. Future work should compare with other approaches, such as recommendation to the entire member population and recommendation of content to produce.

6.2. Directions for Future Research

There is plenty of room for future research. Further studies in other organizations, where social media may be used differently, and experimentation in other periods of the year for longer periods of time, will help reach a better understanding of how to improve recommendations to community owners. In addition, our survey indicated that for some communities, often those focused on a very specific task, recommendations are not likely to be useful. Identifying communities that are suitable for recommendation can also be an interesting future direction.

Another possible direction is to further personalize the recommendations for each community. We saw that factors such as community age and size influence the effectiveness of recommendations and change the way they should be produced. Further personalization can consider the types of items typically shared in the community and their diversity, the number of owners and active members of the community, or the overall pace of activity within the community. For example, for a community that mostly interacts through blogs, recommendation of blog posts may be especially productive.

Our participating owners made many enhancement suggestions that can be further explored in future work, including the ability to post a recommendation to multiple communities owned by the participant; filter out recommendations from sources already associated with most of the community members; explicitly add keywords that would be used for recommendation; add learning capabilities from round to round; use the existing content of the community's blogs, forums, files, and so on. to further enrich the recommendations. To better understand the influence on the community, one participant suggested: "*What about a giveback counter which ticked up every time someone viewed or used a shared document?*"

Although our study was conducted within an enterprise, the technique of recommending content to community owners can also be appropriate for online communities outside the firewall. We are not aware of the existing literature on the topic but hope this study will elicit more work beyond the enterprise scope, which would have to address challenges such as scale, identity management, and motivating factors outside the firewall.

7. CONCLUSIONS

We experimented with recommendation of social media items that enterprise community owners can share with their communities. We found that when it comes to sharing willingness and action, the content of the community, as reflected through its title, description, and tags, plays a central role and is vital for producing recommendations owners would share. This is especially true for smaller and younger communities. We also found that although recommendation is effective across all item types, bookmarks and blogs are the most productive. In addition, presentation of related people and tags contributes to owners' willingness to share items.

Our experiments show a clear and significant effect of the recommendation on community activity over a period of eight weeks. This is reflected in both activity volume and number of contributing members. This effect includes communities in which sharing through the survey was not performed at all and spans members and owners who were not invited to participate. These results indicate that content recommendation to owners can serve as an effective means for boosting participation and contribution in enterprise online communities.

APPENDIX: EMAIL SENT TO COMMUNITY OWNERS

Dear <Community Owner's Name>,

This is an invitation to participate in an experiment, which aims at recommending social media content to owners of IBM Connections communities that may interest the *entire community. As one of the owners of the <Community Name> community, we are interested in your feedback about our recommendations. You will be asked to indicate how likely you are to share IBM Connections items with your community. You will also have the option to instantly share those recommended items with the community.*

The experiment is best viewed in Firefox, Chrome, and IE 9 or above.

Please click here to start. You may be requested to accept our security certificate.

We appreciate your contribution!

Social Technologies Group, IBM Research - Haifa

P.S. You may receive more than one invitation if you are an owner of multiple communities in our sample. Would be great if you can respond for each separately.

REFERENCES

Patrick J. Bateman, Peter H. Gray, and Brian S. Butler. 2011. Research note—The impact of community commitment on participation in online communities. *Inf. Syst. Res.* 22, 4 (2011), 841–854.

Shlomo Berkovsky and Jill Freyne. 2010. Group-based recipe recommendations: Analysis of data aggregation strategies. In *Proceedings of the Fourth ACM Conference on Recommender Systems (RecSys'10)*. ACM, New York, NY, 111–118. DOI:http://dx.doi.org/10.1145/1864708.1864732

Anne Bourhis, Line Dubé, and Réal Jacob. 2005. The success of virtual communities of practice: The leadership factor. *Electron. J. Knowl. Manag.* 3, 1 (2005), 23–34.

Lawrence D. Brown, Tony T. Cai, and Anirban DasGupta. 2001. Interval estimation for a binomial proportion. *Stat. Sci.* 16, 2 (May 2001), 101–117.

Robin Burke. 2002. Hybrid recommender systems: Survey and experiments. *User Model. User-Adapted Interact.* 12, 4 (November 2002), 331–370. DOI:http://dx.doi.org/10.1023/A:1021240730564

Brian Butler, Lee Sproull, Sara Kiesler, and Robert Kraut. 2002. Community effort in online groups: Who does the work and why. *Leadership At A Distance: Research in Technologically Supported Work* (2002), 171–194.

David Carmel, Erel Uziel, Ido Guy, Yosi Mass, and Haggai Roitman. 2012. Folksonomy-based term extraction for word cloud generation. *ACM Trans. Intell. Syst. Technol.* 3, 4, Article 60 (September 2012), 20. DOI:http://dx.doi.org/10.1145/2337542.2337545

Dan Cosley, Dan Frankowski, Loren Terveen, and John Riedl. 2007. SuggestBot: Using intelligent task routing to help people find work in wikipedia. In *Proceedings of the 12th International Conference on Intelligent User Interfaces (IUI'07)*. ACM, New York, NY, 32–41. DOI:http://dx.doi.org/10.1145/1216295.1216309

Marguerite Cronk. 2012. Using gamification to increase student engagement and participation in class discussion. 2012. In *Proceedings of the World Conference on Educational Multimedia, Hypermedia and Telecommunications*. Denver, Colorado, United States, 311–315.

Casey Dugan, Werner Geyer, and David R. Millen. 2010. Lessons learned from blog muse: Audience-based inspiration for bloggers. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI'10).* ACM, New York, NY, 1965–1974. DOI:http://dx.doi.org/10.1145/1753326.1753623

Nader Ale Ebrahim, Shamsuddin Ahmed, and Zahari Taha. 2009. Virtual R&D teams in small and medium enterprises: A literature review. *Sci. Res. Essays* 4, 13 (2009), 1575–1590.

Kate Ehrlich, Michael Muller, Tara Matthews, Ido Guy, and Inbal Ronen. 2014. What motivates members to contribute to enterprise online communities? In *Proceedings of the Companion Publication of the 17th ACM Conference on Computer Supported Cooperative Work & Social Computing (CSCW Companion'14)*. ACM, New York, NY, 149–152. DOI:http://dx.doi.org/10.1145/2556420.2556477

Rosta Farzan, Joan M. DiMicco, David R. Millen, Casey Dugan, Werner Geyer, and Elizabeth A. Brownholtz. 2008. Results from deploying a participation incentive mechanism within the enterprise. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI'08).* ACM, New York, NY, 563–572. DOI:http://dx.doi.org/10.1145/1357054.1357145

Jill Freyne, Michal Jacovi, Ido Guy, and Werner Geyer. 2009. Increasing engagement through early recommender intervention. In *Proceedings of the Third ACM Conference on Recommender Systems (RecSys'09)*. ACM, New York, NY, 85–92. DOI:http://dx.doi.org/10.1145/1639714.1639730

Paul Fugelstad, Patrick Dwyer, Jennifer Filson Moses, John Kim, Cleila Anna Mannino, Loren Terveen, and Mark Snyder. 2012. What makes users rate (share, tag, edit...)?: Predicting patterns of participation in online communities. In *Proceedings of the ACM 2012 Conference on Computer Supported Cooperative Work (CSCW'12)*. ACM, New York, NY, 969–978. DOI:http://dx.doi.org/10.1145/2145204.2145349

David Goldberg, David Nichols, Brian M. Oki, and Douglas Terry. 1992. Using collaborative filtering to weave an information tapestry. *Commun. ACM* 35, 12 (December 1992), 61–70. DOI:10.1145/138859.138867 http://doi.acm.org/10.1145/138859.138867

Bette Gray. 2004. Informal learning in an online community of practice. *J. Distance Educ.* 19, 1 (2004), 20.

Ido Guy, Anat Hashavit, and Yaniv Corem. 2015. Games for crowds: A crowdsourcing game platform for the enterprise. In *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing (CSCW'15)*. ACM, New York, NY, 1860–1871. DOI:http://dx.doi.org/10.1145/2675133.2675189

Ido Guy, Inbal Ronen, Naama Zwerdling, Irena Zuyev-Grabovitch, and Michal Jacovi. 2016. What is your organization 'like'?: A study of liking activity in the enterprise. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems (CHI'16)*. ACM, New York, NY, 3025–3037. DOI:http://dx.doi.org/10.1145/2858036.2858540

Ido Guy, Tal Steier, Maya Barnea, Inbal Ronen, and Tal Daniel. 2012. Swimming against the streamz: Search and analytics over the enterprise activity stream. In *Proceedings of the 21st ACM International Conference on Information and Knowledge Management (CIKM'12)*. ACM, New York, NY, 1587–1591. DOI:http://dx.doi.org/10.1145/2396761.2398478

Ido Guy, Naama Zwerdling, Inbal Ronen, David Carmel, and Erel Uziel. 2010. Social media recommendation based on people and tags. In *Proceedings of the 33rd International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR'10)*. ACM, New York, NY, 194–201. DOI:http://dx.doi.org/10.1145/1835449.1835484

Andrew Hilts and Eric Yu. 2010. Modeling social media support for the elicitation of citizen opinion. In *Proceedings of the International Workshop on Modeling Social Media (MSM'10)*. ACM, New York, NY, Article 3, 4. DOI:http://dx.doi.org/10.1145/1835980.1835983

David Huffaker. 2010. Dimensions of leadership and social influence in online communities. *Hum. Commun. Res.* 36, 4 (2010), 593–617.

Alicia Iriberri and Gondy Leroy. 2009. A life-cycle perspective on online community success. *ACM Comput. Surv.* 41, 2, Article 11 (February 2009), 29. DOI:http://dx.doi.org/10.1145/1459352.1459356

Anthony Jameson. 2004. More than the sum of its members: Challenges for group recommender systems. In *Proceedings of the Working Conference on Advanced Visual Interfaces (AVI'04)*. ACM, New York, NY, 48–54. DOI:http://dx.doi.org/10.1145/989863.989869

Amy Jo Kim. 2000. *Community Building on the Web: Secret Strategies for Successful Online Communities (1st ed.)*. Addison-Wesley Longman Publishing Co., Inc., Boston, MA.

Joon Koh, Young-Gul Kim, Brian Butler, and Gee-Woo Bock. 2007. Encouraging participation in virtual communities. *Commun. ACM* 50, 2 (February 2007), 68–73. DOI:http://dx.doi.org/10.1145/1216016.1216023

Robert Kraut and Paul Resnick. 2011. Encouraging contribution to online communities. 2011. *Building Successful Online Communities: Evidence-Based Social Design* (2011), 21–76.

Ioannis Leftheriotis and Michail N. Giannakos. 2014. Using social media for work: Losing your time or improving your work? *Comput. Hum. Behav.* 31 (2014), 134–142.

Jan Marco Leimeister, Pascal Sidiras, and Helmut Krcmar. 2004. Success factors of virtual communities from the perspective of members and operators: An empirical study. In *Proceedings of the 37th Annual Hawaii International Conference on System Sciences (HICSS'04)*. DOI:10.1109/HICSS.2004.1265459

Pamela J. Ludford, Dan Cosley, Dan Frankowski, and Loren Terveen. 2004. Think different: Increasing online community participation using uniqueness and group dissimilarity. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI'04)*. ACM, New York, NY, 631–638. DOI:http://dx.doi.org/10.1145/985692.985772

Meng Ma and Ritu Agarwal. 2007. Through a glass darkly: Information technology design, identity verification, and knowledge contribution in online communities. *Info. Sys. Res.* 18, 1 (March 2007), 42–67. DOI:10.1287/isre.1070.0113 http://dx.doi.org/10.1287/isre.1070.0113

Christopher D. Manning, Prabhakar Raghavan, and Hinrich Schütze. 2008. *Introduction to Information Retrieval*. Cambridge University Press, New York, NY.

Judith Masthoff. 2004. Group modeling: Selecting a sequence of television items to suit a group of viewers. *User Model. User-Adapted Interact.* 14, 1 (February 2004), 37–85. DOI:http://dx.doi.org/10.1023/B:USER.0000010138.79319.fd

Tara Matthews, Jilin Chen, Steve Whittaker, Aditya Pal, Haiyi Zhu, Hernan Badenes, and Barton Smith. 2014a. Goals and perceived success of online enterprise communities: What is important to leaders & members? In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI'14)*. ACM, New York, NY, 291–300. DOI:http://dx.doi.org/10.1145/2556288.2557201

Tara Matthews, Steve Whittaker, Hernan Badenes, and Barton Smith. 2014b. Beyond end user content to collaborative knowledge mapping: Interrelations among community social tools. In *Proceedings of the 17th ACM Conference on Computer Supported Cooperative Work & Social Computing (CSCW'14)*. ACM, New York, NY, 900–910. DOI:http://dx.doi.org/10.1145/2531602.2531694

Tara Matthews, Steve Whittaker, Hernan Badenes, Barton A. Smith, Michael Muller, Kate Ehrlich, Michelle X. Zhou, and Tessa Lau. 2013. Community insights: Helping community leaders enhance the value of enterprise online communities. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI'13)*. ACM, New York, NY, 513–522. DOI:http://dx.doi.org/10.1145/2470654.2470728

Michael McCandless, Erik Hatcher, and Otis Gospodnetic. 2010. *Lucene in Action, Second Edition: Covers Apache Lucene 3.0*. Manning Publications Co., Greenwich, CT, USA.

Sean M. McNee, John Riedl, and Joseph A. Konstan. 2006. Being accurate is not enough: how accuracy metrics have hurt recommender systems. In *CHI'06 Extended Abstracts on Human Factors in Computing Systems (CHI EA'06)*. ACM, New York, NY, 1097–1101. DOI:http://dx.doi.org/10.1145/1125451.1125659

David R. Millen, Jonathan Feinberg, and Bernard Kerr. 2006. Dogear: Social bookmarking in the enterprise. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI'06)*. ACM, New York, NY, 111–120. DOI:http://dx.doi.org/10.1145/1124772.1124792

David R. Millen, Michael A. Fontaine, and Michael J. Muller. 2002. Understanding the benefit and costs of communities of practice. *Commun. ACM* 45, 4 (April 2002), 69–73. DOI:http://dx.doi.org/10.1145/505248.5052

Michael Muller, Kate Ehrlich, Tara Matthews, Adam Perer, Inbal Ronen, and Ido Guy. 2012. Diversity among enterprise online communities: Collaborating, teaming, and innovating through social media. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI'12)*. ACM, New York, NY, 2815–2824. DOI:http://dx.doi.org/10.1145/2207676.2208685

Mark O'Connor, Dan Cosley, Joseph A. Konstan, and John Riedl. 2001. PolyLens: A recommender system for groups of users. In *Proceedings of the Seventh Conference on European Conference on Computer Supported Cooperative Work (ECSCW'01)*. Kluwer Academic Publishers, Norwell, MA, 199–218.

Michael J. Pazzani and Daniel Billsus. 2007. Content-based recommendation systems. In *The Adaptive Web*. Lecture Notes in Computer Science, Vol. 4321. Springer-Verlag, Berlin, Heidelberg, 325–341.

Adam Perer, Ido Guy, Erel Uziel, Inbal Ronen, and Michal Jacovi. 2011. Visual social network analytics for relationship discovery in the enterprise. In *Proceedings of the IEEE Conference on Visual Analytics Science and Technology (VAST'11)*. Providence, RI, 71–79. DOI:10.1109/VAST.2011.6102443

Jenny Preece and Diane Maloney-Krichmar. 2002. Online communities: focusing on sociability and usability. In *The Human-Computer Interaction Handbook*, L. Erlbaum Associates Inc., Hillsdale, NJ, 596–620.

Jenny Preece, Blair Nonnecke, and Dorine Andrews. 2004. The top five reasons for lurking: Improving community experiences for everyone. *Comput. Hum. Behav.* 20, 2 (2004), 201–223.

Inbal Ronen, Ido Guy, Elad Kravi, and Maya Barnea. 2014. Recommending social media content to community owners. In *Proceedings of the 37th International ACM SIGIR Conference on Research & Development in Information Retrieval (SIGIR'14)*. ACM, New York, NY, 243–252. DOI:http://dx.doi.org/10.1145/2600428.2609596

Inbal Ronen, Elad Shahar, Sigalit Ur, Erel Uziel, Sivan Yogev, Naama Zwerdling, David Carmel, Ido Guy, Nadav Har'el, and Shila Ofek-Koifman. 2009. Social networks and discovery in the enterprise (SaND). In *Proceedings of the 32nd International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR'09)*. ACM, New York, NY, 836–836. DOI:http://dx.doi.org/10.1145/1571941.1572156

Sunanda Sangwan. 2005. Virtual community success: A uses and gratifications perspective. In *Proceedings of the 38th Annual Hawaii International Conference on System Sciences (HICSS'05)*. IEEE Computer Society, Washington, DC, 193c. DOI:http://dx.doi.org/10.1109/HICSS.2005.673

Andrew I. Schein, Alexandrin Popescul, Lyle H. Ungar, and David M. Pennock. 2002. Methods and metrics for cold-start recommendations. In *Proceedings of the 25th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR'02)*. ACM, New York, NY, 253–260. DOI:http://dx.doi.org/10.1145/564376.564421

Lingling Sun and Julita Vassileva. 2006. Social visualization encouraging participation in online communities. In *Proceedings of the 12th International Conference on Groupware: Design, Implementation, and Use (CRIWG'06)*. Springer-Verlag, Berlin, Heidelberg, 349–363. DOI:http://dx.doi.org/10.1007/11853862_28

Etienne Wenger and William M. Snyder. 2002. *Cultivating Communities of Practice: A Guide to Managing Knowledge*. Harvard Business School Press, Boston, MA.

Anbang Xu, Jilin Chen, Tara Matthews, Michael Muller, and Hernan Badenes. 2013. CommunityCompare: Visually comparing communities for online community leaders in the enterprise. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI'13)*. ACM, New York, NY, 523–532. DOI:http://dx.doi.org/10.1145/2470654.2470729

Haiyi Zhu, Jilin Chen, Tara Matthews, Aditya Pal, Hernan Badenes, and Robert E. Kraut. 2014. Selecting an effective niche: An ecological view of the success of online communities. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI'14)*. ACM, New York, NY, 301–310. DOI:http://dx.doi.org/10.1145/2556288.2557348

Received January 2015; revised October 2015; accepted March 2016