Lessons Learned from Blog Muse: Audience-based Inspiration for Bloggers

Casey Dugan, Werner Geyer, David R Millen
IBM T.J. Watson Research
One Rogers Street, Cambridge, MA 02142
{cadugan, werner.geyer, david_r_millen}@us.ibm.com

ABSTRACT
Blogging in the enterprise is increasingly popular and recent research has shown that there are numerous benefits for both individuals and the organization, e.g. developing reputation or sharing knowledge. However, participation is very low, blogs are often abandoned and few users realize those benefits. We have designed and implemented a novel system – called Blog Muse – whose goal is to inspire potential blog writers by connecting them with their audience through a topic-suggestion system. We describe our system design and report results from a 4-week study with 1004 users who installed our tool. Our data indicate that topics requested by users are effective at inspiring bloggers to write and lead to more social interactions around the resulting entries.

Author Keywords
Recommendations, participation, social software, blog.

ACM Classification Keywords
H5.m. Information interfaces and presentation (e.g., HCI): Miscellaneous.

General Terms
Design, Experimentation, Human Factors.

INTRODUCTION
Since the arrival of the first hosted blog tools in 1999, such as livejournal.com or blogger.com, blogging has rapidly gained in popularity. In its 2008 “State of the blogosphere report”, Technorati reported 133 million blog records indexed since 2002 and a rate of 900,000 blog entries created every 24 hours [31]. As they have with many technologies, enterprises have started adopting and molding blogging for professional and workplace use. 12% of all bloggers blog for their company in an official capacity, e.g. to support customer communication or marketing, and 49% of all bloggers blog about their industry, profession, or workplace but not in an official capacity [31]. Companies have also begun to realize the value of blogs as an internal communication tool. Large corporations such as Microsoft, HP, and IBM maintain internal blog sites hosted by their IT department, allowing employees to blog inside the workplace. Employees often start blogging because they expect significant social and informational benefits in their work environment [18] although these expectations are not always met. There is an increasing body of research around the use and utility of blogging in the enterprise. Huh et al. [17] describe how employee blogging supports knowledge sharing and collaboration across internal communities. Jackson et al. [18] report on benefits for individuals and the community, including information sharing, problem solving, getting and giving feedback, engaging in dialog, gaining perspective, gaining company pulse, networking, developing reputation, or building community. Efimova and Grudin [6] found three categories of work-related use: sharing passion for work and direct communication with others, showing a human side of the company, and documenting and organizing work. Not surprisingly, internal employee blogs differ from external blogs in the type of content posted. The majority of external blogs (70%) are largely focused on personal topics [16] whereas employee blogs are mostly work-related or grounded in the work environment [6, 17, 18].

While there are documented benefits of blogging in the workplace, not every employee benefits equally from blogs. Jackson et al. [18] found that less active bloggers get the least value out of blogging. Less than 40% cited work-related benefits whereas more than 70% of heavy users indicated work-related benefits. However, blogging communities are often sustained by only a small core of heavy users [18, 32], suggesting that only few users fully realize those benefits. In our company approximately 3% of all employees have written a blog post, which is similar to numbers reported by Efimova and Grudin [6]; the number of active bloggers is even lower.

Jackson et al. [18] report a number of barriers that prevent wider adoption and use of corporate blogging. Many users do not blog more frequently because they simply have “no time”, which is not surprising in a work environment. Interestingly, heavy users, who benefit most, describe their time as “well spent.” The lack of readership and interesting material “relevant to business goals” is another reason for low participation or users abandoning blogging [18]. Yardi et al. [32] provide further explanation for this lack of readership and commenting through an attention economy model. Another barrier is not knowing how to get started [18], which is consistent with results from our own survey [9], and evidenced by the large number of blogs with no or only few posts (80% of all blogs inside our company have ≤ 5 posts).

The above barriers suggest possible improvements to blogging systems that could potentially help increase activity on blogging sites and thus allow more employees to realize the benefits of blogging in the workplace. In this paper, we propose and study a novel way of engaging users in the blogging process. Two fundamentals of our approach are providing inspirational material for bloggers to write about and connecting bloggers with their audience. By design, blogs are a communication medium which is asymmetric in that there is no obligation for blog readers to engage in social interaction. Interestingly, as mentioned earlier, the role of the reader in blogging is crucial nonetheless. As Nardi et al. [22] state: "*blogs create the audience, but the audience also creates the blog*". Most bloggers desire connecting to an audience and often write with audience attention, feedback, and feelings in mind.

We have designed and built the Blog Muse, a system that capitalizes on the relationship between blog readers and writers. The goal of Blog Muse is to inspire employees to write blog posts by suggesting topics they can write about. In particular, topics can be requested by blog readers, thereby giving the audience a voice and allowing them to share what they would like to read about with the blogging community. Our system then suggests these topics to potential blog writers who can decide whether or not they would like to address the topic requested. The underlying intuition is that users are more likely to blog if they know about their potential audience and the topics of interest. Our system creates a tighter linkage between blog readers and writers and encourages creation of content that matters in a work context by meeting the needs of information consumers.

The design of the Blog Muse and its topic recommendation algorithm was inspired by a survey of potential motivations to blog with 700 employees [9]. In this survey, new and less active bloggers in particular spoke of struggling with coming up with initial topics and desired the "built-in" audience that topic suggestions could provide. Hoping to increase the benefits of blogging for this group, as well as other bloggers, we implemented and deployed Blog Muse to a larger group of users of IBM's internal BlogCentral blogging site in order to better understand its usage in a realistic setting. In particular, we wanted to address the following research questions:

1. How do users respond to the Blog Muse topic recommendations, i.e. how effective is our algorithm and the recommendations it creates at inspiring users?
2. What is the effect of Blog Muse on users and their overall activity on the site, i.e. do users create more blog entries?
3. Do topic recommendations lead to blog entries that are more engaging to the community, in terms of increased interaction or activity, e.g. more commenting, better ratings, or more hits?

This paper reports results from 1004 Blog Muse users over a period of 4 weeks. Our report is structured as follows: We first discuss related work followed by a description of the Blog Muse system and its recommendation algorithm. After describing the design of our study, we present results addressing the aforementioned research questions and we discuss implications for the design of future systems.

RELATED WORK

Recommender systems [25] are increasingly used on social media sites to help users discover content and people of interest. A popular technique to create recommendations is collaborative filtering [15]. This technique makes recommendations based on user similarity either by comparing user ratings or items viewed or purchased. In contrast, classic content-based recommendation approaches utilize the content, e.g. keywords, and match it against a user profile. For example, Pazzani et al. [24] use the words contained on the individual web pages of web sites and built user profiles using websites the user considered "hot". Based on the profile, the system recommends web sites using a Naïve Bayes classifier. Mooney et al. [21] use a similar technique to recommend books. Collaborative and content-based filtering are also often combined into a hybrid system in order to overcome the weaknesses of each approach [3]. For example, Sen et al. [26] built a hybrid alert filtering system using both approaches. More recently, researchers have begun leveraging articulated social network information in order to make recommendations, sometimes referred to as "trust-based" recommendations. Sinha and Swearingen [27] motivate this technique by showing that users prefer recommendations from their friends rather than from online systems. Groh et al. [12] use articulated social network information to generate user neighborhood information. Their experimental data shows that collaborative filtering based on these neighborhoods can improve classic collaborative filtering. Spertus et al. [28] leverage community membership to recommend online communities on the social network site Orkut and they compare several different similarity measures in a large study on the site.

From the above techniques, we recommend topics in Blog Muse using both content-based and network-based techniques for matching topics and users. While typical recommender systems aid in content discovery, Blog Muse focuses primarily on making recommendations to a user for content to create. Similarly, Geyer et al. [11] recommend topics for self-descriptions based on the friend network on an enterprise social networking site. They show that recommendations are effective at increasing the number of profile entries and the number of users with entries. Chen et al. [4] compare content-based and social-network based approaches for recommending people to add to one's friends list. Harper et al. [14] explore personalized invitations to increase participation in an online discussion forum. They show that invitations have an immediate impact in the short term, causing users to write and view more posts. Cosley et al. [5] study four different routing algorithms that assign movies that need editing to MovieLens users and demonstrate that simple but intelligent task routing can increase user contributions. Research that aims at increasing participation and content creation also includes incentive mechanisms. For example, Farzan et al. [8] designed and deployed a point-based incentive system within an enterprise social network site. They find that employees are initially motivated to add new content to the site. Both points and status levels supported by the system contributed to this effect. Sun and Vassileva [29] use motivational social visualization within an online community to increase community awareness and social comparison. They show that this increases the number of content contributions. Compared to the above, Blog Muse targets a blog site and its users and also leverages input from the audience, i.e. blog readers requesting a topic, as a source for recommendations for content users should create. As such, Blog Muse can be considered a socially-inspired recommender system.

The technique of requesting topics in Blog Muse bears similarities to question and answer web sites (Q&A) in which users express their information needs and the community seeks to fill them. For example, Ackerman and McDonald [1] describe an early Q&A system, called Answer Garden, that supports routing of questions to various destinations that could provide an answer, focusing on resolving an unanswered question quickly. Q&A systems have become mainstream social media sites today, see for example, answers.com or answers.yahoo.com. These systems usually feature complex reputation mechanisms in which users can gain points by answering questions correctly, i.e. answers are being rated by the person who posted the question. Blog Muse could be considered a highly contextualized Q&A system that links together questions posed by readers with posts by writers in a blogging system. Rather than awarding points for answers, however, the goal of Blog Muse is to motivate users to write blog entries by offering potential readers and commenters for their blogs.

The problem of routing a topic to a person likely to write about it is also related to expertise location. McDonald [19] leverages social network information for finding knowledgeable colleagues for collaborations in the workplace. Ehrlich et al. [7] describe a system that combines information from email and chat messages with social network information in order to locate experts. McDonald and Ackerman [20] introduce a flexible architecture for expertise location called Expertise Recommender. Their system uses organizationally relevant data to create user profiles. Blog Muse is similar in that it leverages user profile information derived from organizationally relevant data sources such as the corporate directory and internal profile pages.

THE BLOG MUSE APPLICATION

The core concept of the Blog Muse is to close the loop between the people who are looking for content on a blogging site and those who write blog posts. Figure 1 illustrates the social interaction model in a simplified way.

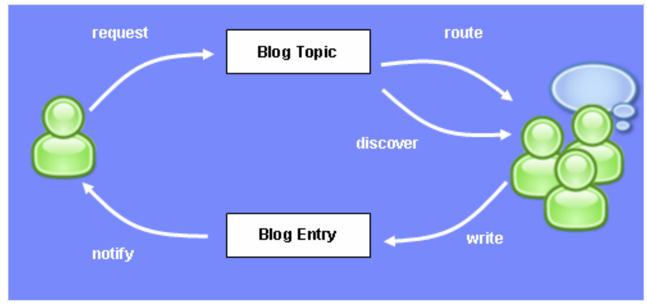

Figure 1. Social Interaction Model.

If users want to read about a certain topic but cannot find any blog posts about it, they can submit a topic request to the Blog Muse. Our system then routes the topic request to users who are likely to write about it, i.e. the system will notify and recommend the requested topic to other users. If a user decides to write a blog post about a requested topic, we notify the requester in return.

The User Interface

Since we were not able to directly modify BlogCentral (the blogging site used within our company) to incorporate the Blog Muse functionality, we gave users the option of installing either or both of two possible add-ons to utilize the Blog Muse: a homepage widget or a Firefox plugin.

The homepage for a suite of social software applications (including BlogCentral) keeps users up to date on activities that are taking place across applications. Additionally, users are able to choose third-party widgets from a catalog and enable them on their homepage. The Blog Muse homepage widget we added to the catalog consists of three primary tabbed views: (Figure 2) "Get Inspired to Write" – where users can see topics recommended to them by the system, (Figure 3) "Ask for a Blog Post" – where users can request topics from the community, and (Figure 4) "Explore Topics and Vote" – where users can see existing topics requests (i.e. recent, popular, or topics from a user's social network).

In the "Get Inspired to Write" view (Figure 2), users can scroll through the blog topics the system has recommended and take the following actions on them: (1) they can write about the topic, (2) they can save a topic and write about it later, (3) they can decline but specify that they would read a blog post about this topic, which adds another vote to the topic and will result in the user being notified if another user later writes about this topic, or (4) they can decline the recommendation. From this view, users may also get back to any recommendations they have saved and take the same actions on those. If a user decides to write about a topic, Blog Muse opens a dialog which allows authoring and publishing a blog post using a rich text editor. The blog post can also be saved to the blogging system as a draft.

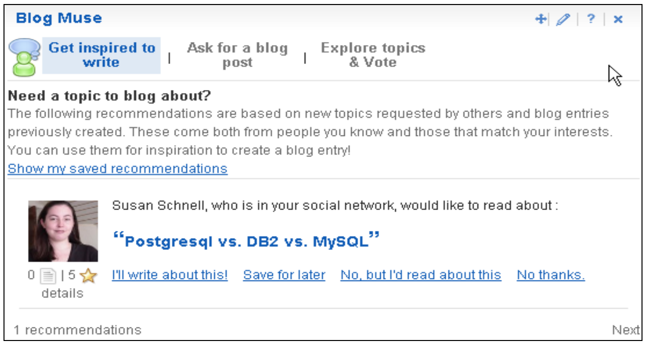

Figure 2. User interface to view recommendations.

In the "Ask for a Blog Post" view (Figure 3), users can request blog posts on a particular topic. Once submitted, users will be notified if entries are later written about their topics. Optionally, requesters can specify the names or email addresses of other users in order to directly request that those users write about this topic. As the user types the topic he/she is requesting, we provide auto-complete functionality in order to reduce the appearance of duplicate topic requests in the system. If the requester selects an existing topic from the auto-complete box, a vote is submitted for that topic. Users can also see topics they previously requested from this view, in order to see how many additional votes those topics have accumulated or to refind an entry written about a topic they requested.

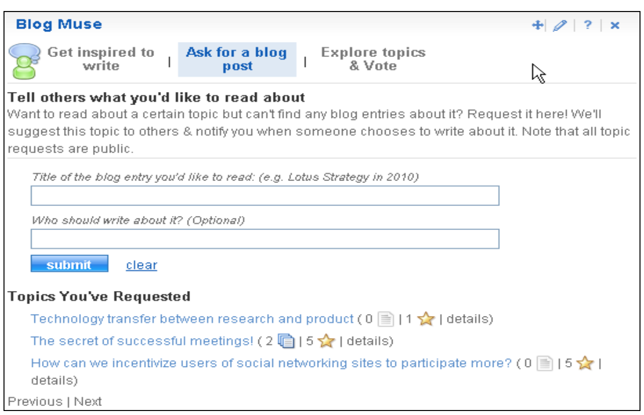

Figure 3. User interface for requesting a new blog topic.

All topics submitted to the Blog Muse are public, and in the "Explore topics & Vote" view (Figure 4), users can serendipitously discover topics to write about through browsing lists showing recent, popular, or topics within a user's social network. Users can also let potential blog writers know that they would be interested in reading about a topic through voting. As with submitting a topic request, when a user votes on a topic, we will notify her when someone blogs about it. Throughout the interface, the number of readers (including the original requester) interested in a topic is shown next to a star icon. Clicking on the "details" link shows the complete list of users who have voted for or requested this topic, as well as any entries written about it. Furthermore, users are given the option of "Recommending to Others." If a user discovers a topic that she thinks someone else should write about or knows about, she can recommend it to them directly by entering their names or email addresses. Finally, a user can also see actual recommendations that were made to her social network. Users are given the option of voting on these if they think they are a good match for their co-workers (e.g. "Vote: I'd like Luis to write this!"). Similar to voting on requested topics, this will increase the "reader" star count on this topic. In addition, we also notify the person in their social network whose recommendation they voted for that the user would like to read about this topic from her.

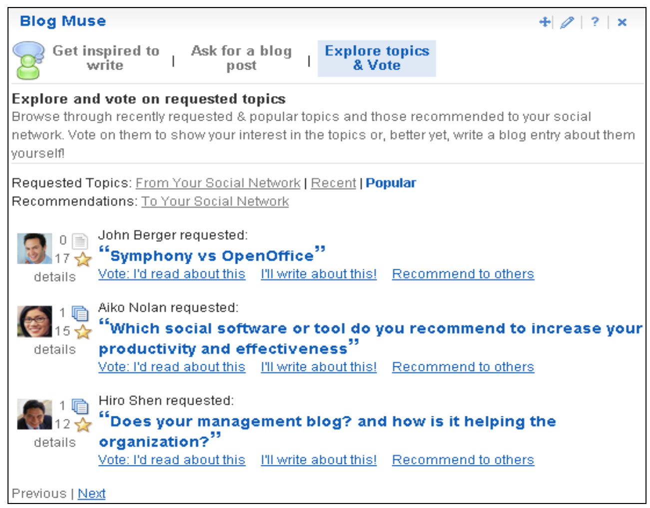

Figure 4. User interface for browsing topics.

In addition to the Blog Muse homepage widget, users could install a Firefox plugin that inserts Blog Muse functionality directly into the user interface of the company blogging site (Figure 5). On the left, a user sees her personalized topic recommendations and can take the same actions on those recommendations as in the Blog Muse homepage widget, including retrieving any saved recommendations. When a user chooses to write about a given recommendation, she is taken to the entry authoring page with the title of the entry pre-populated. When the user publishes a blog post, the Firefox plugin automatically sends this information back to the Blog Muse server. On the right hand side, the user can request topics and specify any users she would like to directly request to write about it. The Firefox plugin does not include the "Explore topics & Vote" view due to space constraints.

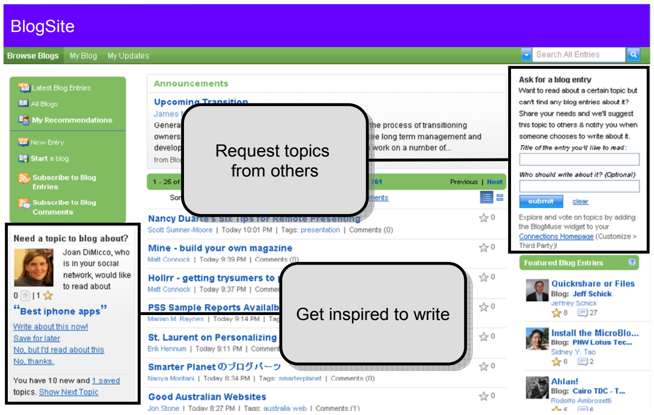

Figure 5. The Firefox plugin augmenting BlogCentral with Blog Muse functionality.

Topic Recommendation Algorithm

As previously described, our goal with Blog Muse is to inspire users to participate by writing blog posts based on personalized topic suggestions made to them. The algorithm used for suggesting and ranking topics to users was derived from a previous survey [9], in which users indicated they would be more likely to blog about topics requested by others than those previously written about. This ranking of recommendations resulting from user requested topics over those topics matched from existing blog entries (also referred to as recycled topics) is shown in Figure 6. User requested topics always get ranked higher. Since the creation of user-requested topics is beyond our control, i.e. dependant on users requesting topics, they get inserted on-the-fly at the top of the recommendation queue of a user whenever a new topic is entered into the system and routed to that user, i.e. if no user-requested topic is available, the user would only see topic recommendations that were pre-computed and matched based on existing topics. We chose this design because we did not want users to have an empty recommendation queue when they first came to the system due to a lack of requested topics.

A direct request, where a set of users were directly specified by the requester at the time of the request, was deemed the most compelling reason in the prior survey, followed by a topic requested by someone in your social network, followed by a requested topic that matches your interests. As such, when a blog reader requests a topic through Blog Muse (as show in Figure 6), the topic is routed to any users directly specified, i.e. *directly routed* ("[user] has requested you write about the following"), plus 50 members of the user's social network, i.e. *network routed* ("[user], who is in your social network, would like to read about the following"), and 50 additional users whose keyword profile matches the topic request using a TF-IDF score and cosine similarity, i.e. *content routed* ("[user] would like to read about the following, which matches your interests/keywords"). Blog Muse then sends out email notifications for user-requested topics that are routed as recommendations to other users. This kind of viral routing aids not only in finding a user who is willing to blog about that topic but also leads to additional users installing our application.

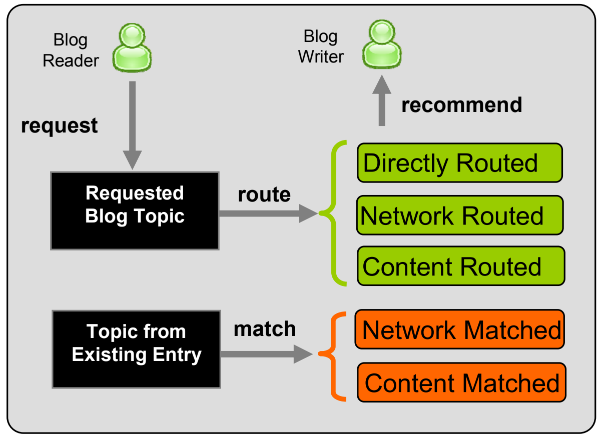

Figure 6. The topic recommendation process model

For a given user, the existing topics (blog titles "recycled" from existing entries) recommended to her were either previously written about by her social network, i.e. *network matched* ("[user] who is in your social network, wrote about the following") or matched her keyword profile using a TF-IDF score and cosine similarity, i.e. *content-matched* ("[user] wrote about the following, which matches your interests/keywords"); these were shown to users in alternating order.

RESEARCH METHODS

Blog Muse was deployed as an add-on to an IBM internal blogging site (BlogCentral) which has more than 70,000 registered users (15,269 of which have authored entries) and 17,129 blogs (including both personal and group blogs). The site has been active since November, 2003. We conducted a controlled experiment of the Blog Muse system between May 9th, 2009 and June 11th, 2009. We assigned BlogCentral users to three different groups: **control**, **baseline**, and **experimental**. Each group consisted of 1001 users and was a random stratified sample from the seven categories of blogging activity levels as defined in [9], ranging from users who had no BlogCentral activity, users with a blog but no entries, users who just commented, occasional users, to very active bloggers. Note that the overwhelming majority are not active or even occasional bloggers. For each user, we also required being able to make at least 4 existing-topic network-matched recommendations and 4 existing-topic content-matched recommendations. Users from all three groups received an email at the beginning of the experimental period, asking them to participate in a research trial and giving them instructions for installing both Blog Muse interfaces. The control group also had access to the tool, but did not receive personalized recommendations and was not able to discover requested topics like the other two groups. Their interface consisted of only the "Ask for a Blog Post" tab (Figure 3). The baseline and experimental groups had a similar user experience, in which the "Get inspired to write" tab was open by default. However, users in the baseline group received recommendations randomly selected from previously written about topics[1] whereas the experimental group received personalized recommendations using the Blog Muse algorithm described in the previous section. The experimental group also received periodic email notifications when they were directly requested to write about a topic by another user or they indirectly received a topic requested by another user (i.e. by network or content routed). Notifications of direct requests to write about a topic were suppressed for both the baseline and control groups during the experimental period.

A system such as Blog Muse is very much dependent on contributions from users enabling it to work. For example, users must first request blog topics in order for us to effectively compare human-requested topics to previously written about topics. In order to jumpstart the creation of user requested topics, we sent emails to 3,003 random additional users ("contributors") asking them to participate by installing the tool and submitting requests for blog topics. Their interface was similar to the experimental groups, though the "Ask for a blog post" tab was open by default, and they also received personalized recommendations. In addition, users who came upon the tool by other means, for example, through the viral nature of our routing algorithm or by discovering it through the widget catalog on the homepage, and did not fall into any of the previously described groups, were given complete functionality.

RESULTS

During the experimental period, a total of 509 users downloaded and installed at least one of the two available Blog Muse tools: 155 in the control, 158 in the baseline, and 196 in the experimental group. 168 users (67 from baseline and 101 from experimental group) responded to 2,587 recommendations. A total of 225 topics were requested by users. These requests came from those in our three groups plus additional users who installed the tools through the viral effect of the routing algorithm, through word of mouth, discovery on our internal web site, and through the emails that were sent to help bootstrap the system with user-requested topics. A total of 495 additional users who did not fall into our three groups installed at least one of the two tools.

The following analysis is structured by the three research questions outlined earlier in the paper. Unless otherwise stated, our analysis is based on the 509 users in the three groups.

[1] The explanation given for random recommendation was the same as in the content-matched case for the experimental group.

Response to the Blog Muse recommendation algorithm

Random versus Blog Muse recommendations

The design of the Blog Muse recommendation algorithm was informed by a recent survey [9] as described in the previous section. In order to understand how well this algorithm works, we compared acceptance of recommendations ("good" recommendations) between the baseline (random) and the experimental group. A recommendation was considered *good* if users responded in any of the following three ways: "I'll write about this" "Save for later", or "No, but I'll read about this." Figure 7 shows that the Blog Muse algorithm has a significantly higher rate of "good" recommendations compared to our random baseline (64.0% versus 34.6%, $F[1,166] = 28.113$, $p<0.001$). While the superior results of the recommendations based on topic similarity and social network relatedness over a random blog topic are not very surprising, we wanted to confirm that the algorithm used provided our anticipated boost in performance.

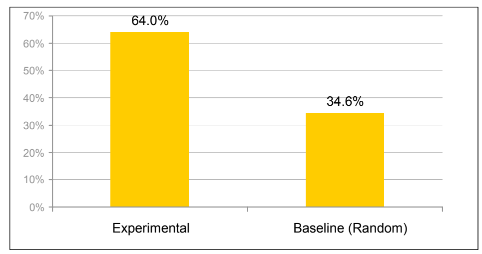

Figure 7. Acceptance of recommendations for the Blog Muse algorithm versus random existing-topic recommendations.

Requested Topics versus Recycled Topics

The Blog Muse algorithm recommends topics requested by users (requested topics) as well as topics from previously written blog posts (existing topics). A within-subject analysis of recommendations from user-requested topics versus existing / recycled topics shows that topics requested by users are significantly more accepted than existing topics with 56% acceptance versus 37.5% as shown in Figure 8 ($F[1,80]=16.033$, $p<0.001$). For the above analysis we considered only users who had at least one requested and one existing topic.

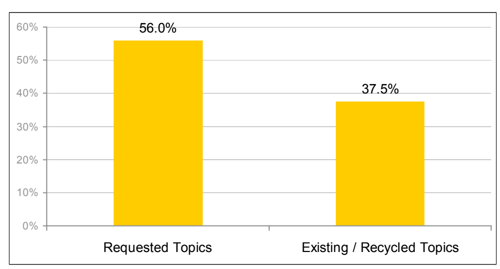

Figure 8. Acceptance of recommendations from user requested topics versus existing topics.

A more detailed analysis of requested versus existing topics by routing algorithm and computation method is shown in Figure 9.

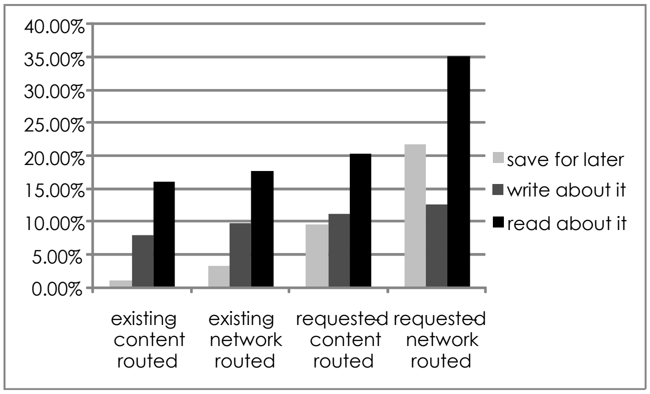

Figure 9. User response to user-requested topics and existing topics broken down by routing and matching method.

Of the responses to the topics that were rated favorably: 142 were "write about it," 330 were "save for later," and 647 were "read about it." Note that recommendations from user-requested topics that were directly routed to individuals (i.e. a requester specified a set of recipients) are not shown due to very few being created. The data reveals that user-requested topics routed to the network are significantly more likely to result in a user writing about a topic, followed by requested topics routed to users whose profile matches the topic. Recommendations based on existing topics both from the social network (network matched) and from a profile match (content matched) have very low acceptance rates and seem unsuitable as recommendations of topics to write about (χ^2 (9, N = 3571) = 312.0, p< 0.001).

Impact of Voting

The above data shows that user-requested topics with a built-in audience, as determined by a user requesting it, are more likely to be written about. When we designed Blog Muse, we wanted to reinforce the notion of an audience by allowing other users to vote on topics, thus growing the audience beyond a single user. To understand the impact of this feature on "write" responses to our recommendations, we analyzed all topics that had votes and compared them to topics without votes. Figure 10 shows that topics with votes were more than 6 times as likely to be written about than topics without votes (χ^2 (1, N = 225)= 9.651, p<0.01).

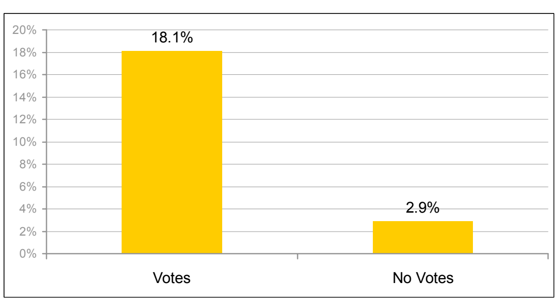

Figure 10. User response to topics with and without votes.

Effect of the Blog Muse on User Activity

Control Group versus Experimental

We compared users from each of our three groups to determine the effect of Blog Muse on their overall activity on the site during the experimental period. A higher percentage of users in the experimental group (29.08%) added at least one personal blog entry during the experimental period compared to those in the control group (26.45%) as shown in Table 1. However, we were not able to find a significant difference in the average number of personal blog entries added by the control group (1.58) and the experimental group (1.44) during the period[2] (F[2,136]=117, p >.05).

Table 1. Comparison of new blog entries

	CTRL	BASE	EXP
# users who downloaded Blog Muse	155	158[3]	196
Total # new blog entries created	245	175	282
Average # new blog entries per user	1.58	1.11	1.44
% who added at least one entry during period (of those who installed tool)	26.45%	26.58%	29.08%
# new entries created from Blog Muse recommendations		1	12

Engagement Around Blog Muse Entries

While we did not observe differences in the overall quantity of blog entries created across the experimental groups, we considered that there may have been differences in the blog entries created as a result of a Blog Muse recommendation. To explore this, we further examined the 13 blog entries created by the study participants (Table 1). We also included in this exploratory analysis, an additional 28 blog entries that were created from Blog Muse recommendations by other users who were not in our three groups but had learned about Blog Muse and had installed the application. Of these 41 new blog entries, 33 entries were the result of a new topic specifically requested by a Blog Muse user, while 8 entries were created from "recycled" existing topics.

Comments, Ratings, and Hits on Blog Muse Entries

During the experimental period, 2,865 entries were authored on BlogCentral that were not created through the Blog Muse system. We compared the 41 Blog Muse entries to all others created on the site, looking for differences in comments, ratings, and hits. A summary of these results can be seen in Table 2.

Of the 2,865 entries not created through Blog Muse, 18.78% received comments, compared to 43.90% of the

---

[2] Since the number of blog entries follows a power law distribution, we used a log transformation in our analysis. We report the raw units in Table 1 for readability.

[3] Three users were removed from the baseline group due to manager requests.

Blog Muse entries. The average number of comments on Blog Muse entries (0.85) was higher than the average number of comments on the other BlogCentral entries (0.43) ($F[1,2904] = 10.9$, $p < 0.001$)[4]

Table 2. Average number of actions on blog entries

Mean	ratings	comments	views
Blog Muse (n=41)	0.49	0.85	23.46
All others on BlogCentral (n=2865)	0.16	0.43	19.46
User-Requested (n=33)	.61	1.00	26.45
Existing / Recycled (n=8)	0	.25	11.13

On BlogCentral, users are able to give entries that they would recommend to others a single "star" rating. Of the entries created through Blog Muse, 29.27% received ratings compared to only 11.31% of those not created through Blog Muse. The average number of stars on Blog Muse entries (.49) was significantly higher than the average number on the other entries created (.16) ($F[1,2904]=16.6$, $p < .001$).

The Blog Muse entries were also viewed more often on average (23.46) than those not created through Blog Muse (19.46) ($F[1,2904]=7.96$, $p < 0.005$).

Of considerable interest is that when looking at only the subset of the Blog Muse entries that came from topic requests made by users (as opposed to the recycled/existing topics), the average number of comments, ratings, and views increases even more compared to BlogCentral. As can be seen in Table 2, entries from user-requested topics received almost four times as many ratings on average (.61 compared to .16) than all others created on BlogCentral and twice as many comments on average (1.00 compared to .43) They also received more views (26.45) on average than all others created on BlogCentral.

DISCUSSION

We expected that our trial would confirm that the Blog Muse recommendation algorithm performs significantly better than a random baseline algorithm. More importantly though, topics requested by users are valued higher than recycled existing topics and, given the responses from our users, seem to be the only feasible way of inspiring both current and new bloggers to write. This data confirms findings from a previous survey [9] in which user-requested topics were ranked higher than existing topics.

However, the survey also indicated that recycled topics could potentially provide inspirational material to writers. This was one of the reasons why we also included recycled topics in our Blog Muse recommendation algorithm. Another reason why we decided to include them was to be able to prepopulate users' topic recommendation queues. Note that the creation of user-requested topics was not fully under our control and we felt it was important that users who visit the system the first time have some number of recommendations. However, the results from our deployment provide limited support for this. Given the high acceptance of user-requested topics over existing topics, we have revisited that decision. Recommender systems in particular can suffer from users losing trust in the system and the content that is being recommended after a number of bad suggestions [30]. We may have potentially turned a number of users, who would have been happy to write about the right requested topic, off the system by recommending too many of these poor, existing topics. The same could be said about the users who were in our baseline group and received random recommendations. Both of these could have contributed to the application having lower adoption than expected – users are unlikely to tell their coworkers about a cool new tool that is not working well for them.

The data from our trial does show that blog entries created with our system engage the blogging community by receiving more interaction and traffic, i.e. we observed significantly more ratings, comments, and hits than other entries created on BlogCentral during the experimental period, in particular for user-requested topics. It appears that creating a tighter linkage between blog writers and their audience is indeed beneficial. One possible reason for why we did not observe a significant increase in the quantity created could simply be a displacement of writing activity, i.e. entries created through Blog Muse replaced entries that would have been created otherwise. These targeted entries might even represent time better spent by the employee as they provide clear value to a known subset of the community. This kind of targeted writing may be particularly useful in bringing less active bloggers the kind of perceived benefits active bloggers report while not increasing their writing volume. But more research is necessary to determine if authors perceive additional benefits from blogging by using Blog Muse and to compare the effects on the different categories of bloggers.

While Blog Muse provided blog readers and authors with the benefit of more targeted blog topics, there are also potential costs associated with such an application. There is an opportunity cost associated with increasing the quantity of blogging: time spent by employees authoring or reading entries is time not spent on other work-related tasks. Indeed, Blog Muse actively seeks out more employees to spend time writing or reading entries requested by their colleagues. Increasing the number of blog entries within an organization also contributes to the growing information overload problem. We believe that the benefits of more targeted blog entries outweighs these costs, but additional research is needed to better understand how Blog Muse actually changes the cost/benefit structure of the corporate blogging economy.

[4] As above, the analysis of ratings, comments and views was done after a log transformation of the variables. Raw units are reported.

While we know from prior survey data that direct requests are the most positive reason for choosing to write a blog post on a given topic, direct recipients for a requested topic were rarely specified in Blog Muse in practice. Part of that may have to do with users being uncomfortable directly asking others for a blog post about a certain topic. This hints at the social pressure implications of the topic requests in our system (such as "what if my boss asks me to write an entry?") and warrants further study. Despite this, the system could also reduce the overhead placed on the user in making this decision, for example by suggesting people to directly route the request to. Beyond this, topic requesters could also be brought more fully into the routing process, which is transparent to end-users today. For example, the system could present to the user those people it believes most likely to write about the topic (through both social network and content matches) and allow the user to select a subset of these to actually route the request to. This would have benefits in potentially increasing the accuracy of our routing algorithm as well as "upgrading" the reason why a given user is being routed a particular request (e.g. from "[user], who is in your social network, would like to read about" to "[user] would like **you** to write about the following").

In addition, the social features in Blog Muse could be pushed even further. For example, users might be able to leave comments with their votes or ask follow-up questions regarding a requested topic – neither of which is supported by the system today. Clarification of requests, such as a detailed description in addition to topic title, is a particularly compelling angle to pursue, as topics are currently limited to 255 characters – or the maximum length of the blog title on BlogCentral.

Our data shows that the key idea of creating a closer connection between authors and their audience has merits, as demonstrated by the impact of votes on authors choosing which topics to blog about and the increased audience interaction around Blog Muse entries in the form of comments, ratings, and hits. This idea can be generalized beyond blogging to other forms of social media such as bookmarks, videos, photos, etc. In such a system, users would be able to tell others what they would like to be able to see or read about alongside the content that the sites' other users are currently contributing. Many sites today highlight popular content or trending topics from their contributors, which can grow into world-wide memes – and the same could be done for audience requests. For example, it is not hard to imagine a rush amongst YouTube users to be the first to upload a video in response to a spike in users looking for a replay of a controversial awards show moment. Depending on the social media type, our system might also work more or less effectively. Some social media types are easier to share compared to writing a blog entry (see e.g. [10]) and this might lead to a higher number of contributions created through user-requests.

In Blog Muse, requests from users were manually submitted. But the request process could become semi- or fully automated if integrated with search functionality, e.g. every time a user performs a search on a site and finds no matching content, a request could be submitted or if the user indicates that the search results did not contain what he/she was looking for. As the number of blog entries created is ultimately bounded by the number of topics requested, such an automated approach might have also led to an increase in the quantity of entries produced.

CONCLUSION

We presented Blog Muse, a novel audience-based topic suggestion system for blog readers and writers. The goal of our system is to inspire bloggers to write and to make the entries created more engaging to the blog community. Our system capitalizes on the importance of an audience to blog writers by connecting them with potential readers. It also provides an opportunity for readers to participate and find content (through topic requests) that matters to them in a work context. The importance of the audience in the blogging ecosystem has only recently been studied in more depth [2, 32]. In Nardi et al.'s [23] study on why people blog they nicely state how for one of their subjects the "*community became his muse; [and] his poems developed as a conversation between himself and other bloggers*." Our research contributes to and complements previous research on the value of the audience in blogging.

The lessons learned from our user study validate our design and provide useful input for us for the next release of Blog Muse. While we were able to show that user-based topic suggestions can be beneficial for a blogging site since they lead to more interactive blog entries driving more traffic, more work needs to get a deeper understanding of the overall implications of topic suggestions on the adoption and use of enterprise blogging. To that end, we are currently planning on deploying Blog Muse site-wide for a longitudinal study that will shed more light on how the overall growth of the site is affected and how our system will be able to alter or remove participation barriers for readers and writers. We hope to see an overall increase in readership of less popular blogs. Ideally, this increase will be caused not only by shifting attention from more popular to less popular/known blogs but also by increasing overall adoption of the site. A longitudinal study would also provide an opportunity to measure how our system affects bloggers with different activity levels, to better understand the overall cost and benefits, whether or not it increases reader satisfaction (i.e. produces more meaningful content for enterprise users), and if it reinforces social benefits by creating new ties between employees.

REFERENCES

1. Ackerman, M.S. and McDonald, D. W. 1996. Answer Garden 2: Merging Organizational Memory with Collaborative Help. In *Proc CSCW'96.*

2. Baumer, E., Sueyoshi, M., Tomlinson, B. 2008. Exploring the role of the reader in the activity of blogging. In *Proc. CHI'08*, 1111-1120
3. Burke, R. Hybrid Recommender Systems: Survey and Experiments. *User Modeling and User-Adapted Interaction*. 12(4), 331-370.
4. Chen, J., Geyer, W., Dugan, C., Muller, M., and Guy, I. 2009. Make new friends, but keep the old: recommending people on social networking sites. In *Proc CHI '09*. ACM, New York, NY, 201-210.
5. Cosley, D., Frankowski, D., Terveen, L., and Riedl, J. 2006. Using Intelligent Task Routing and Contribution Review to Help Communities Build Artifacts of Lasting Value. In *Proc CHI'06*. 1037-1046
6. Efimova, L., Grudin, J. Cross Boundaries: A Case Study of Employee Blogging. In *Proc HICSS'07*, IEEE Press.
7. Ehrlich, K., Lin, C., and Griffiths-Fisher, V. 2007. Searching for experts in the enterprise: combining text and social network analysis. *Proc Group'07*, 117-126.
8. Farzan, R., DiMicco, J., Millen, D. R., Dugan, C., Geyer, W. and Brownholtz, E. Results from deploying a participation incentive mechanism within the enterprise. *Proc. CHI'08*, (2008), 563-572.
9. Geyer, W., Dugan, C. 2009. Inspired by the Audience – A Topic Suggestion System for Blog Writers and Readers. To appear: *Proc. CSCW'10*.
10. Geyer, W., Dugan, C., DiMicco, J., Millen, D. R., Brownholtz, B., Muller, M. 2008. Use and reuse of shared lists as a social content type. In *Proc CHI'08.*
11. Geyer, W., Dugan, C., Millen, D.R., Muller, M., Freyne, J. 2008. Recommending Topics for Self-Descriptions in Online User Profiles. In *Proc RecSys'08*, 59-66.
12. Groh, G., & Ehmig, C. 2007. Recommendations in Taste Related Domains: Collaborative Filtering vs. Social Filtering. In *Proc ACM Group'07*. 127-136.
13. Gurzick, D., Lutters, W. G. From Personal to the Profound: Understanding the Blog Life Cycle. In *Proc CHI'06*, 827-832.
14. Harper, F.M., Frankowski, D., Drenner, S., Ren, Y.,Kiesler, S.Terveen, L.,Kraut, R., and Riedl, J. Talk Amongst Yourselves: Inviting Users to Participate in Online Conversations. In *Proc. IUI'07*, (2007), 62–71.
15. Herlocker , J. L., Konstan, J.A. , Riedl, J. 2000. Explaining collaborative filtering recommendations. In *Proc ACM CSCW'00*. 241-250.
16. Herring, S.C., Scheidt, L.A., Bonus, S., Wright, E. Bridging the Gap: A Genre Analysis of Weblogs. In *Proc HICSS'04*, IEEE Press.
17. Huh, J., Jones, L., Erickson, T., Kellogg, W. A., Bellamy, R. K., and Thomas, J. C. BlogCentral: the role of internal blogs at work. In *Proc CHI '07,* 2447-2452.
18. Jackson, A., Yates, J., Orlikowski, W. Corporate Blogging: Building community through persistent digital talk. In *Proc HICSS'07*, IEEE Press.
19. McDonald, D. W. 2003. Recommending collaboration with social networks: a comparative evaluation. In *Proc ACM CHI'03*, 593-600.
20. McDonald, David W. and Ackerman, Mark S. 2000: Expertise Recommender: A Flexible Recommendation System and Architecture. In *Proc CSCW'00*, 231-240.
21. Mooney, R. J., & Roy, L. 2000. Content-based book recommending using learning for text categorization. In *Proc ACM DL'00*.195–204.
22. Nardi, B. A., Schiano, D. J., and Gumbrecht, M. Blogging as social activity, or, would you let 900 million people read your diary?. In *Proc. CSCW '04*.
23. Nardi, B. A., Schiano, D. J., and Gumbrecht, M., Swartz, L. 2004. Why we blog. *CACM*, Vol. 47, No.12, 41-46.
24. Pazzani, M. J., Muramatsu, J., & Billsus, D. 1996. Syskill webert: Identifying interesting web sites. *AAAI/IAAI*, Vol. 1, 54-61.
25. Resnick, Paul and Varian, Hal. Recommender Systems, introduction to special section of Communications of the ACM, March 1997, vol. 40(3).
26. Sen, S., Geyer, W., Muller, M., Moore, M., Brownholtz, B., Wilcox, E., & Millen, D.R. 2006. FeedMe: a collaborative alert filtering system. In *Proc. ACM CSCW'06*. 89-98.
27. Sinha, R. and Swearingen, K. Comparing Recommendations made by Online Systems and Friends. *Proc. DELOS-NSF Workshop on Personalization and Recommender Systems in Digital Libraries*, (2001).
28. Spertus, E., Sahami, M., and Buyukkokten, O. 2005 Evaluating similarity measures: a large-scale study in the Orkut social network. In *Proc SIGKDD'05*. 678-684.
29. Sun, L. and Vassileva, J. 2006. Social Visualization Encouraging Participation in Online Communities. *Groupware: Design, Implementation, and Use*, 349—363.
30. Swearingen, K. and Sinha, R. 2002. Interaction Design for Recommender Systems. In *Proc DIS'02*.
31. http://technorati.com/blogging/state-of-the-blogosphere/
32. Yardi, S., Golder, S. A., and Brzozowski, M. J. Blogging at work and the corporate attention economy. In *Proc CHI '09*. 2071-2080.