Enhancing Directed Content Sharing on the Web

Michael S. Bernstein, Adam Marcus, David R. Karger, and Robert C. Miller
MIT CSAIL
32 Vassar Street, Cambridge MA
{msbernst, marcua, karger, rcm}@csail.mit.edu

ABSTRACT
To find interesting, personally relevant web content, people rely on friends and colleagues to pass links along as they encounter them. In this paper, we study and augment link-sharing via e-mail, the most popular means of sharing web content today. Armed with survey data indicating that active sharers of novel web content are often those that actively seek it out, we developed FeedMe, a plug-in for Google Reader that makes directed sharing of content a more salient part of the user experience. FeedMe recommends friends who may be interested in seeing content that the user is viewing, provides information on what the recipient has seen and how many emails they have received recently, and gives recipients the opportunity to provide lightweight feedback when they appreciate shared content. FeedMe introduces a novel design space within mixed-initiative social recommenders: friends who know the user voluntarily vet the material on the user's behalf. We performed a two-week field experiment (N=60) and found that FeedMe made it easier and more enjoyable to share content that recipients appreciated and would not have found otherwise.

Author Keywords
Social link sharing, blogs, RSS, friendsourcing

ACM Classification Keywords
H5.2. Information interfaces and presentation (e.g., HCI): Miscellaneous.

General Terms
Design, Human Factors

INTRODUCTION

In the struggle to manage information overload on the web, we might characterize two extreme groups: those who drink from the firehose, and those who carefully sip from the stream of content. The firehose-drinkers consume immense amounts of web content to find as much interesting material as possible. They use aggregators and tools such as RSS (Really Simple Syndication) to aid their search. Those who sip in small doses prefer to trust a small set of sources, reading relatively little and missing interesting gems that do not cross their path. Neither strategy is perfect; both sides express interest in seeing more interesting content.

CHI 2010, April 10–15, 2010, Atlanta, Georgia, USA.

To help the firehose-drinkers and the sippers both get more of what they want, our work builds on social interactions that help power information awareness today: people share web pages by e-mail, by talking in person, and by posting to social networks [10, 17]. We are in pursuit of a socially translucent design: one reflecting and empowering these existing social patterns [11]. Social link sharing is often high-quality and personalized: quality is vetted by people you trust, and personalization is implicit when your social network uses its notion of your interests to forward you URLs. However, the process is inhibited by sharers' fear of spamming friends and forwarding old or irrelevant material.

Our goal in this work is to *understand* the social processes behind web content sharing and to *support* those processes by introducing a novel tool to facilitate such sharing.

To understand the process, we investigate the process of social link sharing, building on earlier survey research (e.g., [10, 17]). We find that e-mail is still the dominant sharing medium despite the proliferation of social sharing tools, that topic interest is the biggest determiner of recipient enjoyment, and that a small number of recipients typically signals the most relevant content. The strongest predictor of interest in sharing content is an interest in seeking out new content, rather than measures of a person's social capital.

Armed with the knowledge that aggressive content consumers are also the most prolific sharers, we designed a tool to support consumers in directed sharing of web content with those who want to receive more but do not want to drink directly from the firehose. The tool, FeedMe (feedme.csail.mit.edu), is a plug-in for the RSS reader Google Reader (www.google.com/reader). FeedMe provides two primary mechanisms to support sharing: *Recommendations* and *Social Awareness and Feedback*. FeedMe learns recipients' content preferences based on previously shared content, and suggests potential recipients inline with RSS posts being viewed. Recommendations reduce the amount of effort required to share to two clicks: one click to select a recommended recipient, and one more to send. This approach draws research on recommender systems (e.g., [19, 22, 26]) into a social sharing tool. In parallel, social awareness helps sharers avoid spamming by making visible information such as number of shared items.

FeedMe introduces a novel design space within mixed-initiative social systems: the user mediates recommendations not for themselves, but on behalf of someone they know. Most mixed-initiative systems ask the user to vet

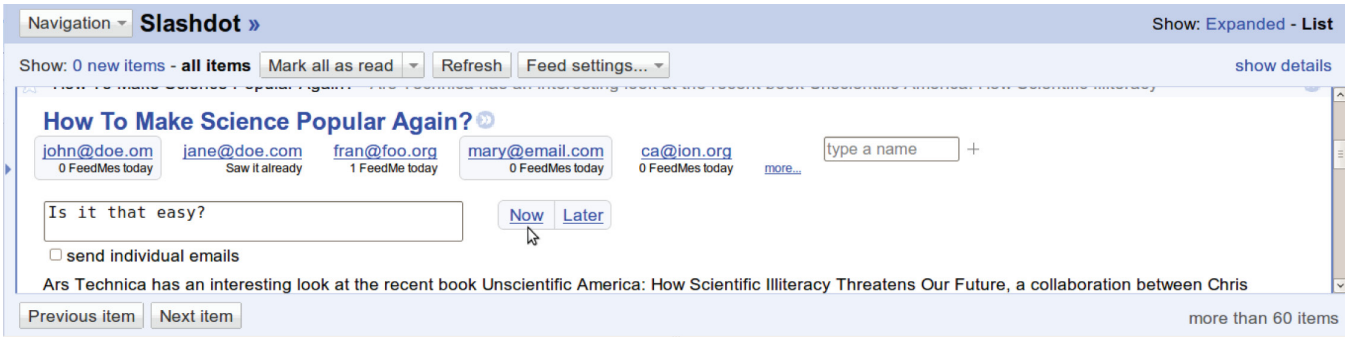

Figure 1. The FeedMe plug-in for Google Reader suggests friends, family, and colleagues who might be interested in seeing the post that you are reading. This user has selected john@doe.com and mary@email.com out of the list of 5 recommendations. The "Now" button sends an e-mail immediately; the "Later" button queues the item in a digest of multiple messages.

suggestions for their own use, like a search autosuggest helping to formulate queries. Instead, FeedMe's users act as gatekeepers *for someone else*. This approach addresses two challenges with recommender systems: they require training, and recommendation mistakes result in time wasted reading unwanted content. We shift the training burden from receivers to sharers, who incur little cost because they have already read the article. Further, whereas traditional recommender systems' mistakes cause the receiver to waste time reading an article, sharers can quickly gloss over faulty recommendations in FeedMe.

In this paper we contribute: 1) an investigation of social link sharing practice: what motivates it, and what prevents it from happening more; 2) a system supporting the directed sharing of links with contacts through recommendations, social awareness, and social feedback; 3) a mixed-initiative social interaction mechanism that minimizes false positives by using friends' knowledge of each other; and 4) a two-week field experiment demonstrating that FeedMe makes link sharing easier for sharers and benefits recipients.

RELATED WORK

We first examine the practice of information sharing on the web. Individuals who are the most successful sharers become knowledge brokers in their local network, known variously as Ehrlich and Cash's information mediators [8], Paepcke's contact brokers [21] and Allen's technological gatekeepers [2]. We directly pursue these information mediators as a user group. Erdelez and Rioux [10] found that the web was the most common source of encountered information that their study population shared with others, and that in-person conversation and e-mail were the most popular means of information transfer. Marshall and Bly built a taxonomy of information sharing: sharing to educate or raise consciousness, sharing using common interests to raise rapport, and sharing to demonstrate knowledge of the recipient's unique interests [17]. Our survey work affirms and extends their research in an era of social sharing tools.

We contribute to a growing set of literature on blog reading. Baumer found blog reading to be a relaxing habit, and in fact that many blog readers do not feel that they suffer from information overload [4]. We extend this discussion by studying the role of blog reader as sharer. Baumer and Fisher developed the Smarter Blogroll, which uses topic analysis to portray trending topics [3]. Like Smarter Blogroll, we perform text analytics to aid the user experience. Other blog-reading interfaces include BLEWS, a visualization of discussion in the political blogosphere [12], and NusEye, which focuses on term co-occurrences [7]. Rather than visual analytics, as in this trio of tools, FeedMe focuses on the blog sharing experience.

FeedMe's sharing recommendations build on previous work in recommender systems. Montaner et al. provide a taxonomy of Internet-based recommendation systems [20]: popular techniques include collaborative filtering (e.g., GroupLens [22]) and mining browsing history (e.g., WebWatcher [15]). FeedMe bears closer resemblance to the Do-I-Care Agent, which explored how collaboration technologies can support recommender systems by allowing agents to communicate [1]. Rather than focusing on discovery as recommender systems do, FeedMe focuses on sharing. To generate sharing recommendations, we must build profiles for users who are neither searching for content nor contributing to their profile. We accomplish this goal by engaging friends who have already read each post. Since FeedMe recommends *people*, we also draw on work in expertise recommenders to match people to a piece of information (e.g., [19]). In characterizing social matching [26], Terveen and McDonald note that profile data can either be provided by the user or mined from a social network, and explore the privacy, trust, and reputation implications of such systems. FeedMe takes these notions to one extreme: it builds on existing social connections, and requires no profile to start.

FeedMe uses humans to filter the content that a recipient accesses. PHOAKS by Terveen et al. [27] took this approach, determining popular web content by crawling Usenet for frequent mentions of a given webpage. This method has manifested itself more recently in services such as in Google News (news.google.com) and Digg (www.digg.com), but these approaches are not personalized. Both services offer personalization through collaborative filtering or sub-communities, but the user must be actively involved to benefit. They also do not have the final human verification that FeedMe introduces: an algorithm still makes the final decision on what to promote.

What tools do you use to share content?	
E-mail	38
Talking in person	29
Social network sites (e.g., Facebook, MySpace)	18
Chat or instant message	16
Twitter	7
Blogging platforms (e.g., Blogger, Wordpress, Tumblr)	6
News aggregators	4
Social bookmarking sites (e.g., delicious)	3
StumbleUpon	1
RSS/Feed Reader (e.g., Google Reader)	1

Table 1. E-mail and talking in person are by far the most common means of sharing today. (N=40, multiple responses allowed)

How do you go about finding and viewing new web content?	
Regularly visiting favorite web sites	36
E-mails from people I know	34
Posted links from friends on social network sites	15
E-mail summaries and digests that I signed up for	14
Google Reader	13
Twitter	7
News aggregators (e.g., Digg, Reddit)	7
Other RSS/Feed Reader	5
StumbleUpon	3
Social bookmarking services (e.g., del.icio.us)	3
Tumblr	1

Table 2. Web sites and e-mails are the most common means of finding new web content (N=40, multiple responses allowed).

Finally, we nod to other work which uses the FeedMe system name or paper title [6, 25].

SURVEYING EXISTING PRACTICE

Sharing web content is woven into the fabric of web citizens' social lives. To better understand the process, we extend previous survey research of web sharing habits [10, 17]. To follow, we present two surveys that inform the design of FeedMe. These findings helped us explore the design space – a more rigorous and in-depth study would be a noteworthy research undertaking on its own.

How Does Sharing Happen?

To begin, we were interested in the social issues that moderate sharing, the tools that individuals use to share, and the kinds of sharing activity that are appreciated or disliked by recipients. We proceeded via an online survey.

Our survey investigated both the sharing and the receiving of web content. We inquired whether receiving links from friends and family is a positive experience, whether participants would be interested in receiving more links than they do now, and which qualities make for good and bad shares. We investigated whether forwards sent to groups had a different quality than those sent only to the individual, and which factors motivate a reply to a link-share. We asked similar questions from the sharer's perspective: What are the significant motivators and fears when deciding whether to forward a link? Who do people share with most, and why? The survey consisted of a mix of multiple choice questions, free-response questions, and Likert scales.

Which is the strongest motivator when you share links?	
I know the person would appreciate hearing about it	37
I like being seen as a source of interesting web content	2
I'm looking to comment or start a conversation about it	0
Sharing a link makes it more likely that we can find it later	0

Table 3. Sharing is strongly motivated by a sense that the receiver would be interested in what you're looking at. (N=39)

Which is the biggest concern you have when you share links?	
I'm not sure whether the link is relevant enough	13
They might have seen it already	7
It's too much effort for me to send the links	6
I have sent the recipient(s) too many links recently	5
It's awkward to contact someone out of the blue	4
I'm not sure that the contents are of high enough quality	3

Table 4. Being unsure of relevance to the recipient's interests is the largest concern cited with sharing. (N=38)

We recruited 40 participants from Amazon Mechanical Turk. Participants were paid $0.20 for their participation – a fairly large amount on Mechanical Turk, which is dominated by tasks for less than $0.10. Mechanical Turk demographics are in line with our desired user group: generally college-educated, 58% female, 20-40 year old Americans [18]. We based our survey on best practices for Mechanical Turk user studies, such as by making it difficult to answer our questions dishonestly via free-response [16]. We saw little evidence of responders cheating the system; many wrote in-depth responses, and we kept all responses after inspecting the written text.

Results

Email is the Dominant Sharing Medium. For sharing and receiving links, we found e-mail to be the most common route (Tables 1-2). Despite the proliferation of social sharing tools like Facebook and Twitter, e-mail is still preferred for its ubiquity and consistency. Everyone has an e-mail address and most people check their email constantly. "I am too busy for the other forms," one participant noted, "I check e-mail throughout the day." These results extend the findings of Erdelez and Rioux from ten years ago [10]. 75% of our participants also reported participating in face-to-face link sharing: the topic would come up in conversation, and one person would show the link while both are present.

Topic Interest Drives Enjoyment. Participants articulated two categories of URLs that they enjoy receiving: topics of interest and entertainment. Topics of interest are highly individual: they range from finance to politics, Michael Jackson ("because I am a great fan"), and educational technology. These categories are often specific; as one participant reported, "Those who know my politics usually send me very pointed articles – no junk." Entertainment links largely consisted of humorous articles and YouTube videos.

These same categories also generated the most ire: missing the mark was the most commonly cited reason for disliking

a share. Participants cited politics and YouTube as sources of irritation as often as they were cited as source of enjoyment. Of links on politics or religion, one participant reported, "Don't try to conform me." Of YouTube videos, another said, "I could care less about a cat boxing."

Sharers are largely aware of recipients' goals and the potential pitfalls of sharing. Of 39 respondents, 37 stated that their strongest motivator for sharing was the knowledge that the receiver would appreciate it (Table 3). In parallel, sharers' largest concern was determining whether the link would be relevant enough to the person or group (Table 4).

Link Sharing Is Burdensome When It Is a Repetitive Firehose. A sharer's failure to rate-limit their posts is a commonly cited frustration. One participant discussed a particular individual who sent them 10-20 items per day, "blindly forwarded on." "They send me Fwd:Fwd:Fwd: type emails," another complained of an annoying sender. Receivers disapprove of old news – things they have seen before. This aversion poses a challenge to sharers, who have incomplete knowledge of what the recipient has seen. This situation was the second-most common concern sharers reported (Table 4). Rate-limiting was also common: "I don't want to take a chance of annoying someone."

Small Audiences Are Best. The fewer the number of people receiving a link, the more interest recipients have in reading it. In general, links shared with smaller groups are more targeted to the individual's interests. Participants described that they are more likely to read and respond to links sent only to them. "I don't click on links from mailing lists," one participant admitted. When sharing, participants reported sending the links to small numbers of people.

Friends Are the Most Common Target. Links are typically sent to close friends. Sharers articulated two reasons for this: they are more certain of close friends' tastes, and they are already in regular communication.

Recipients Want More. We asked respondents to rate on a Likert scale the statement, "If guaranteed to be links I'd like, I would be interested in receiving more links from people I know than I do now." We found that recipients are willing to receive more: the median response was 6 out of 7 (μ=5.3, σ=1.3). This result suggests an opportunity for users to share more if we can motivate them to do so.

Who Are the Active Sharers?

Where our first survey investigated the dynamics of sharing, our second sought to uncover the characteristics of the most active sharers. We were interested in two hypotheses from our own experience and previous work [10, 17]: that *sharers are especially social individuals*, and that *sharers seek out and experience large quantities of web content*.

To investigate the social orientation of our participants, we turned to previous work on *social capital*. Social capital broadly refers to the resources (or "capital") that develop through maintenance of social relationships. Individuals or communities can turn to those with whom they have established social capital for favors or support. In our work we focus on two specific types of social capital utilized by Ellison et al.: *bridging social capital* and *bonding social capital* [9]. We adapted scales consisting of ten Likert-scale questions from Ellison et al. to measure bridging social capital and bonding social capital.

Seeking scale
I spend a large amount of free time viewing web content.
I am rarely one of the first people to know about interesting web content. (reversed)
I follow many sources of web content for updates.
I check for new or updated web content very often.
I often seek out updates on topics relevant to my interests or my job using the internet.
I often seek out entertaining posts, jokes, comics and videos using the internet.
I often seek out updates on people or groups I know using the internet.
I read or skim the titles of all the posts made to my favorite web sites or blogs.
I rely on tools that aggregate popular web content from many sources: for example, Google News, Google Reader or Digg.
I rarely rely on the internet for content relevant to my interests. (reversed)

Table 5. Questions in the Seeking scale, investigating interest in finding and consuming web content.

Sharing scale
People I know see me as a source of interesting or funny web content.
When I see something I like on the internet, my first thought is often, "Who else would enjoy seeing this?"
My friends tend to share more web content than I do. (reversed)
I often post interesting web content to public places like my IM status, my Facebook profile, or Digg.
I often send interesting web content to people I know or to groups that I belong to.
I often send a link to someone I know after I am reminded of it during a conversation.
Sharing links is a way I keep in touch with people I know.
I often tell people I know about my favorite web sites to follow.
I rarely share links with people I know. (reversed)
I often talk about the web content I have seen with other people.

Table 6. Questions in the Sharing scale, investigating interest in passing web content on to others.

Bridging social capital focuses on the aspects of social capital corresponding to weak ties – loose connections of individuals who see each other occasionally, such as might be found in the reaches of an enterprise or a neighborhood. For example, one bridging social capital scale item was: "I come in contact with new people all the time." High bridging social capital is associated with many of the benefits of the classic "Strength of Weak Ties" article by Granovetter [13], including access to information and jobs. Bonding social capital is generated between strong ties, or close groups of family or friends. An example bonding scale item: "There is someone I can turn to for advice about making very important decisions." We investigate these constructs because they can teach us about what kinds of relationships are important to the information gatekeepers in link-sharing scenarios.

We then constructed two scales of our own: a *seeking* scale and a *sharing* scale (Tables 5-6). Our seeking scale measures how much time and interest an individual invests in finding interesting or entertaining web content. The sharing scale measures how likely an individual is to share web content with friends, family and colleagues. Both scales consist of ten Likert-scale questions iteratively developed and refined via pilot studies. These scales have the weak-

Predictor of sharing scale	β	t	p-value
Seeking scale	.74	8.38	< .001
Bridging social capital scale	.22	2.36	< .05
Bonding social capital scale	.01	0.14	.33

Table 7. Interest in seeking is strongly correlated with interest in sharing, much more so than bridging social capital or bonding social capital. (Adj. R^2=0.56)

ness that they do not measure exactly how often our participants engage in such activities, instead focusing on expressed interest. Pilot studies found self-report to be unreliable for 'how often' questions, so we utilized the vaguer wording in the scales to capture overall sentiment.

We presented the forty questions in random order for each participant. This survey was distributed via Mechanical Turk to a group of 100 individuals. Participants were again paid $0.20. We looked for anomalies in the distribution of Likert scores, and kept 99 survey results after dropping one participant who responded 'neutral' to all questions.

Results

An individual's score on the seeking, sharing, bridging social capital and bonding social capital scales is the average of their answers on all ten 7-point Likert scales. For all scales, Cronbach's alpha (a measure of agreement) was good: between .7 and .9. We verified that all scales were distributed normally and did not exhibit heteroscedasticity problems. Following Ellison et al. [9], we performed principal components factor analysis with varimax rotation to verify that factors loaded on the correct constructs, and found reasonable but not perfect correspondence. These statistical results indicated that we satisfied the mathematical assumptions necessary for a regression analysis, and that the questions that we devised were in fact testing four different underlying concepts.

We then performed an ordinary least squares regression using seeking, bridging social capital and bonding social capital as independent variables and the sharing score as a dependent variable (Table 7). We found that interest in seeking is strongly correlated with interest in sharing (β=.74, $p < .001$), explaining more than half the variance in sharing scores. The social capital measures explain only 3% more variance when added to the model.

We note that bridging social capital is more strongly correlated with sharing than bonding social capital. So, expressed interest in maintaining a large number of weak ties impacts the amount of sharing that is done. However, our survey results indicate that most sharing happens with strong ties and not weak ties. Further research is needed to explain this result, especially since social capital measures can be correlated. We hypothesize that while those with large weak tie networks do have interest in sharing information, when sharing, people feel most comfortable communicating with close friends.

Most importantly, this result indicates that interest in seeking out new web content is much more important than measures of sociality in determining how much users share web content with friends, family and colleagues. Thus, we believe that tools to encourage information sharing can profitably focus on information seekers.

Survey Limitations

We intended our survey to inform design choices and to extend the inquiry of the general phenomenon. Being a survey, however, it is limited to self-report. A survey can lead participants to report more socially desirable answers than might be true – for example, we note a general trend in our data toward being above neutral on all scales. Our sample may also be biased due to our use of Mechanical Turk for data collection. In particular, Turkers may not represent a completely accurate cross-section of Internet users. While the data may not be considered conclusive for hypothesis-testing, the effects are relatively strong and are sufficient for inspiring design ideas. In addition, previous work studying Mechanical Turk demographics suggests that N=100 will reach a wide-ranging group [18].

FEEDME

Our investigations suggested that we could increase the amount of productive, personalized sharing by: 1) targeting users of RSS feed readers, who have demonstrated interest in seeking out web content; 2) addressing concerns of whether a recipient might be interested in a post, reducing the effort barrier to sharing; and 3) mitigating social concerns associated with the sharing process, like spamming and worrying whether a recipient already saw the link.

FeedMe is a plug-in for Google Reader that suggests contacts who might be interested in seeing the content currently being viewed (Figure 1), and provides social awareness and feedback mechanisms to ease spamming concerns. To follow, we describe FeedMe's two major components: sharing recommendations, and social awareness and feedback.

Recommendation Interface

FeedMe injects a recommendation interface under the title of every post viewed in Google Reader (Figure 1). The recommendation interface suggests individuals with possible interest in the post being viewed. The recommendations make sharing a two-click process: click to confirm the recipient, then click the "Now" button to send an e-mail. Users can optionally add a comment that will be prepended to the e-mail. If multiple receivers are selected, the e-mail goes to all of them; the user also has the option to send separate e-mails rather than cc'ing each recipient.

If interested, the user can display more recommended recipients by clicking "more" to reveal another row of recommendations. If the desired contact has not been recommended or if the user has not shared with the contact before, the user can enter the contact's e-mail address in an autocompleting textbox. This box is populated with the user's Google contacts. When the user first uses FeedMe, no recommendations are available and the user must bootstrap using autocomplete. As the user shares, the system recommends past recipients for new posts.

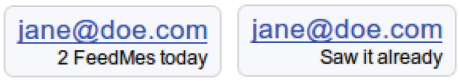

Figure 2. Load indicators reflect the number of items sent today (left) and whether the receiver has seen the post already (right).

The recipients do not need to be FeedMe users, use an RSS reader, or invest effort in profile authoring. This imbalance is desirable because the majority of recipients do not use an RSS reader, and thus would never use FeedMe. However, many sharers live within an RSS ecosystem. So to train our recommender, we utilize the efforts of the (relatively fewer) FeedMe users: FeedMe models a recipient's interest by tracking the posts shared with that recipient.

Social Awareness Information and Social Feedback

FeedMe's social features are intended to display useful information about the receiver to the sharer, give the sharer more control over how the link is sent, and give the receiver a lightweight feedback mechanism.

Load Indicators

Our survey participants expressed fear about spamming contacts. To help the user gauge the likelihood of being perceived as spammy, FeedMe provides social awareness information with its recommendations (Figure 2). A primary concern is whether the recipient has seen the item already, so FeedMe displays "Seen it already" if the recipient has received the link from another FeedMe user or if the recipient is a FeedMe user and viewed the item in Google Reader. This alert depends on the information that FeedMe can observe, such as FeedMe shares and Google Reader viewership. The interface also helps the sharer gauge how overwhelmed the recipient is by counting FeedMe e-mails from FeedMe since midnight. For example, if the recipient has received 2 FeedMe e-mails from one user and 3 from another, the interface displays "5 FeedMes today."

Digest E-mails

If sharers are worried about sending too many e-mails, they can opt to click "Later" instead of "Now" when sending the e-mail (Figure 1). "Later" queues the message into a digest e-mail that is sent out to recipients twice a week when there are pending shared items. A sharer can queue as many items as desired, knowing that only one e-mail will be sent.

One-Click Thanks

Replying to e-mails enables conversation, but recipients may want to express appreciation for the shared post without writing a detailed response. To facilitate this, FeedMe provides a lightweight thanking mechanism to let the sharer know when a recipient appreciates the content. If Dan Olsen were to share a post, a link with the action text "Send Dan Olsen a One-Click Thanks!" is added to the e-mail below the post title. When a recipient clicks the link, he or she is taken to a confirmation page with a thanks leaderboard. The leaderboard counts the number of times each of the sharer's recipients has thanked the sharer, inspired by social games like Collabio [5]. Simultaneously, the sharer is notified of the thanks by e-mail.

IMPLEMENTATION

We implemented the user interface for FeedMe as a Greasemonkey script. Greasemonkey is a plug-in for the Firefox web browser that facilitates the modification of a web site's code and interface. DOM listeners determine when the user has shifted their attention to a new post. For each post, FeedMe sends an AJAX request for recommendations. The server is implemented using the Django framework and stores data in a MySQL database.

FeedMe constructs a recommendation profile for each user who has received a shared post. To do this, it builds a bag of words model for each recipient composed of words that have appeared in posts previously recommended to them. The algorithm concatenates post title, feed title and content of every post sent to the recipient, then tokenizes the result, performs word stemming, and removes common stop words. Words are weighted by term frequency-inverse document frequency (TF-IDF) [24], so that popular words in posts sent to the recipient are more salient.

The recommendation algorithm uses the standard Rocchio approach, computing cosine distances to each friend of the sharer to the post and ranking the friends' distances [23]. The server creates a TF-IDF word vector for the post, then compares that vector to the vector representing each recipient the sharer has shared with in the past.

EVALUATION

To evaluate FeedMe's impact on sharing habits, we performed a two-week field experiment. We recruited 60 participants via blogs and e-mail lists who were regular users of Google Reader and Firefox. We paid participants $30 for two weeks of Google Reader use with FeedMe installed.

Participants filled out a pre-study survey containing our 10-item seeking and sharing scales. Median participant age range was 26-30, and 46 were male. Many participants were students; others included consultants, designers, an editor, an entrepreneur, a music teacher, a theater technician and a patent agent. The mean seeking and sharing indices for the participants were 5.7 and 4.84, respectively, and t-tests confirm our expectation that both seeking and sharing indices were higher than the general Internet users we surveyed earlier: $t(158) = -6.375$, $p < .001$ and $t(158) = -3.215$, $p < .01$. Participants also shared 30-day usage statistics that Google Reader makes available before they began using FeedMe. The median participant read 1,598 posts from 52 feeds in the month preceding the study, shared 0 posts from Google Reader using the built-in e-mail interface (though many sent more, max. 224) and publicly shared 5 posts.

Field Experiment Design

FeedMe takes two approaches to facilitate sharing: recommending potential recipients and social awareness and feedback. We designed a study to understand whether these features are useful and how they impact sharing, in a 2 (recommendations) x 2 (social) design. All factors were fully balanced and randomized.

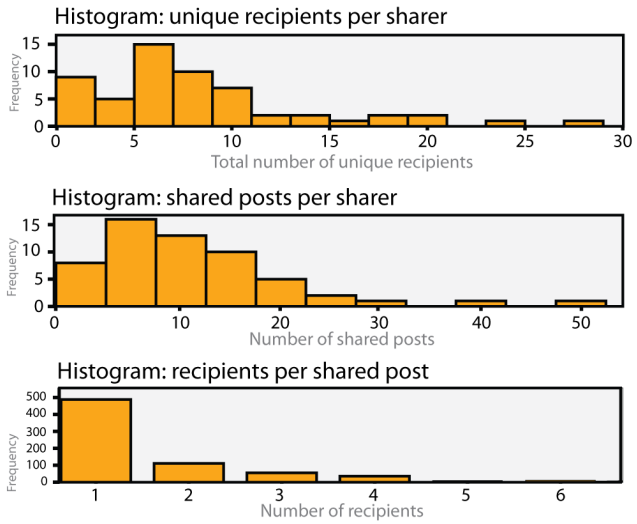

Figure 3. Typically, users shared with small numbers of individuals and addressed each message to one recipient.

Recommendations were either fully enabled or not shown – in either condition, the user could also use an autocomplete textbox to manually add an e-mail address. This factor was within-subjects: participants tried each interface for a week, half receiving recommendations only in the first week, and half receiving them only in the second week. We did not add a second control group with random recommendations: we wished to focus on the social impact of sharing rather than the specific algorithm, and piloting had shown the Rocchio algorithm good enough for eliciting this feedback.

Social features were either fully enabled or fully disabled for the length of the study. Disabling the social features removed information about number of messages received today, whether the recipient had seen or received the link already, the ability to digest e-mails for later, and the ability of recipients to send One-Click Thanks. The social factor was between-subjects, so participants remained in their group for the entire study. We chose to make social features a between-subjects variable to simplify the user experience: four (2x2) configurations would be more difficult for participants to remember and compare.

Halfway through the study and again at the end of the study, we asked participants to complete a survey about their experience. The survey asked Likert scale and free response questions about that week's interface, including ease of sharing and concern about spamminess.

Results

In the results to follow, we report that both sharers and receivers found real benefit in FeedMe. Receivers reported that 80% of shared posts were novel content, and that they were glad to receive the posts. Fully 31% of shared posts had at least one One-Click Thanks. Sharers also enjoyed the tool: 18 participants continued to use the tool a week after the study ended. Participants told us that recommendations made sharing easier and were significantly in favor of it compared to the control interface. Load indicators put

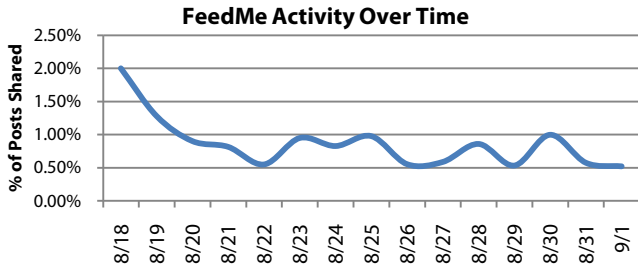

Figure 4. After the initial rush of activity, participants continued to use FeedMe to send a consistent percentage of posts viewed.

sharers at ease and digests freed some users to send many more posts than other study participants.

Usage Trends

Of the 60 users who were initially enrolled in the study, 58 used FeedMe until the end of the two weeks and responded to all of our survey questions. These participants shared a total of 713 items using FeedMe, 0.84% of the 84,667 posts viewed while FeedMe was enabled in Google Reader. The median number of viewed posts during the period, normalized out to 30 days, was 1,639 – roughly in line with reading trends prior to the study (median 1,598). Figure 3 shows three histograms of usage statistics: unique recipients, shared posts, and recipients per post. There is a right skew to all three distributions: 81% of our users shared with 10 or fewer recipients, most participants shared 20 or fewer posts, and most posts were shared with a single recipient.

It is tempting to argue that 20 shared posts in two weeks is a low figure, and that participants tried and then discarded FeedMe. Sharers were, however, consistently using the tool. The first two days saw higher activity levels, after which sharers shared a relatively constant number of posts per viewed article through the two weeks (Figure 4). We required participants to have the tool installed, but we did not require them to share – the uniformity of sharing across the study suggests that users did not lose interest. As further evidence, two days after the end of the study, 25 of the 60 participants were still using FeedMe to share posts; a week after the end of the study, 18 participants were still using FeedMe. This evidence is indirect, but we consider the voluntary continued usage to be implicit positive feedback. However, given the relatively small number of shared posts, we proceeded with our summative evaluation largely via qualitative assessments, augmented with usage statistics.

All versions of FeedMe had a large effect on the amount of sharing occurring within the Google Reader interface. A paired t-test comparing the number of posts that sharers e-mailed using Google Reader in the 30 days before the trial (μ=2.7, σ=.86) to the number of posts that sharers e-mailed using FeedMe (extrapolated from 14 days to 30; μ=26.5, σ=20.9) is highly significant: $t(57) = 8.447$, $p < .001$. This data is of course not convincing by itself due to the Hawthorne effect, but it suggests that we successfully transitioned information seekers to sharers.

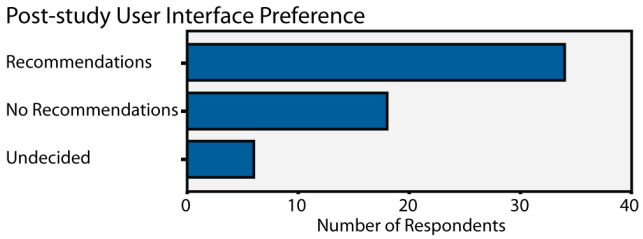

Figure 5. Participants reported a significant preference for the recommendation interface ($p < .05$).

To begin to understand FeedMe's impact, we need to investigate those most impacted by the software. Arguably, this group is not the sharers, but the larger number of receivers who had an unexpected windfall of web links.

Receiver Feedback

Receivers' impressions of FeedMe are an important primary benchmark of success. We emailed everyone who had received at least one FeedMe shared post with a short survey, offering entry in a $30 raffle in compensation. The survey randomly selected up to five posts that the recipient had received via FeedMe. For each post, we asked 1) whether the recipient had seen the link somewhere other than the FeedMe e-mail, and 2) how glad the receiver was to have received that post, on a 7-point Likert scale.

We received responses for 166 shared posts on behalf of 64 receivers. We found that receivers were generally glad to have received the information: the mean Likert response was 5.1 (σ=1.6). Receivers also indicated that the vast majority (80.4%) of posts were only encountered through FeedMe. Since the posts were generally enjoyable, it is clear that FeedMe then directly benefited the recipients, who saw more than they would have otherwise.

We conclude that recipients did not feel spammed by FeedMe, were pleased by the shared posts, and were more up-to-date thanks to the novel posts shared by their friends.

Recommendation Interface

Participants viewed the recommendations as a useful means of lowering the effort barrier to sharing. When asked about their favorite part of FeedMe, participants often mentioned the recommendations. One participant appreciated the "keyboard-free, convenient emailing of articles to friends I share with all the time (and have therefore built up a record of in FeedMe)." The recommendations appeared to achieve FeedMe's design goal of accelerated sharing. "I can rapidly click names of people I regularly contact," a participant shared; another reported his favorite feature to be "the speed with which you can share content (without any new tabs or pages)." Participants who preferred the no-recommendation interface did so for reasons of clutter and waste of vertical pixels in Google Reader. FeedMe's recommendations were also occasionally off-target, especially with individuals e-mailed only once.

We asked users to express a preference for either the version of FeedMe that contained recommendations or the one that did not. Using a practice described by Hearst [14], we named the interfaces "Aspen" and "Sierra" for comparison purposes. Two researchers coded the freeform responses as favoring recommendations, favoring no recommendations, or undecided (Figure 5). The codings agreed at a .938 level as measured by Cohen's kappa, indicating almost perfect correspondence. A third party arbitrated disagreements. A chi-square test indicates a clear preference for the recommendation interface ($\chi^2(1, N=58) = 4.92$, $p < .05$), with nearly twice as many participants preferring recommendations to no recommendations (34 to 18).

Social Awareness and Feedback

Demand for the social features was high – participants who spent the two weeks without social features (re-)invented them in feedback surveys. Nine of the 30 users with social features mentioned digests, activity statistics, or One-Click Thanks as being their favorite feature in FeedMe. "I could worry less about annoying [my friends]," one participant described. When asked what feature of FeedMe would make them feel more comfortable sharing more, 14 of 28 users without social awareness and feedback indicated that knowledge of how overloaded recipient are would help them feel more comfortable sharing, whereas only 3 of the 30 users with social features made such a claim. The difference between these two groups is significant, as verified by a Chi-Square test with Yates' correction ($\chi^2(1, N = 58) = 9.34$, $p < .01$). Thus, we believe that the social features went far to address awareness concerns.

Receivers and sharers both appreciated the One-Click Thanks feature. Of 349 shared posts sent in the social-enabled condition, 108 (30.9%) received at least one thanks. An informal sampling of four Facebook feeds revealed that a similar percentage (~30%) of posts receive at least one Like — an equal engagement from a much larger audience. One recipient who contacted the researchers expressed that One-Click Thanks made it simple to express gratitude for messages which they previously felt pressure to provide an in-depth response to and would typically not respond to at all. The thanks leaderboard did not stimulate competition, but it had the benefit of making user activity visible, thus providing social proof of FeedMe usage.

The Seen It Already indicator was not triggered often because our sharers had largely distinct sets of friends. The feature's usefulness would presumably be improved as entire social circles adopt FeedMe. One participant reported: "I feel like the 'saw it already' feature could be a sleeper hit for me, it doesn't seem special at first but could be really spectacular to know who has seen or shared an item already." Feedback suggested that it would be particularly useful when sharing with other feed reader users.

Opportunities for Improvement

The clearest concern with FeedMe is related to the choice of e-mail for delivering messages. Some users considered email to be sacred and professional. One shared: "I'm pretty conservative about invading people's email space…I worry that they will take 'real' email from me less seriously" if

they also receive lighter, comedic content such as cartoons. The perceived problem is that e-mail is a push medium: recipients are forced to look at the links along with more important information. "Email is a more direct way to communicate," one participant explained, "and I feel that articles that are I read are more like 'ambient' information." For this reason, some power users preferred media they could firehose, such as the public sharing option on Google Reader. Only 5 out of 38 respondents to our original survey indicated that this kind of rate-limiting was their most pressing concern, but it was clearly a theme of the FeedMe feedback. We can think of two explanations: 1) active information seekers are more sensitive to e-mail crowding than average Internet users; 2) FeedMe addressed other concerns successfully enough to make rate-limiting the most pressing remaining concern.

Limitations of the Study

In order to participate in our study, participants had to be Google Reader users with the latest version of Firefox and the ability to install Greasemonkey. Participants who fit this profile are likely to be power users, biasing the kind of users on whom we base our conclusions. Such users often had established norms for sharing with friends, such as mailing lists or IRC channels, and were potentially more sensitive to increasing e-mail traffic to recipients. These biases might have resulted in less sharing in situations where the general population of users might not be so sensitive or have outlets other than email on which to share interesting content.

DISCUSSION AND FUTURE WORK

We return to the notion of FeedMe as a novel design for a mixed-initiative social recommender system. Instead of marshalling machine learning in the service of information filtering, FeedMe marshals it in the service of information sharing. FeedMe recommends content to intermediate sharers who can efficiently and effectively verify that the receiver would be interested. One benefit of this approach is that users may be more tolerant of errors when acting on another person's behalf than when the recommendations are for themselves. We found that FeedMe users appreciated recommendations when they were accurate, and generally did not mind if they were wrong. This design has low marginal cost for sharers – when they have already taken the time to find interesting content, filtering requires relatively little additional overhead. FeedMe opens up a design space with potential impact on other domains where AI is still error-prone, for example expert finding: "We think that your friend Sanjay can answer this question about Nikon cameras: [question]. Is he a good person to ask?"

Another benefit of FeedMe's approach is that it can build personalized models without the recipients' participation. We can rely on active RSS readers to install the tool and build the models, because it aids sharing, but the recipients need not do anything to benefit. As evidence of this positive imbalance, only 6.2% of the shared posts in our study were sent to another FeedMe user. We have no traditional means to model interest for the recipients of the remaining 93.8% of shared posts, so FeedMe's passive learning through sharing is important in the vast majority of cases. However, there is room for recipient involvement: though we expected to work with no receiver feedback, we found that recipients readily adopted lightweight social cues such as the One-Click Thanks to signal preferences. We thus plan on augmenting FeedMe to better-inform senders which receivers appreciated previous shares, and to improve the recommendation algorithm with recipient feedback.

We expect that FeedMe's recommendation algorithm could be improved. For example, integrating a notion of how often a sharer communicates with a recipient would likely improve our accuracy. This approach would promote common recipients and suppress recommendation of one-offs. Our survey also suggested that forwarding patterns differ based on whether the content is informative or entertaining; we might also build a simple 'entertainment classifier' and provide other kinds of social feedback in this instance, such as how viral the content is.

FeedMe's field experiment highlighted issues with today's sharing media. Though participants preferred a recommendation system for ease-of-use, neither the recommendations nor the social feedback had a strong impact on the amount of sharing that occurred. Primary among sharers' concerns was an aversion to spam: participants are hesitant to share too much via a non-ignorable feed such as e-mail. Unfortunately, there is no low-priority queue for receivers as pervasive as e-mail. Users wanted other means of sharing more in line with their individual practice: IRC, IM, or RSS. However, IRC and RSS do not have general adoption, and many view IM as a high-priority feed. Social network streams may provide an attractive alternative.

Privacy issues are worth addressing briefly. FeedMe combines information in posts sent by all sharers to model its receivers, but this has privacy implications if a friend is recommended an article on a sensitive topic. For example, a user reading posts about a medical condition may be surprised to find a recommendation suggesting that a friend is a good match, inadvertently alerting the user that the friend has a medical condition. A simple solution would be to build separate models for each sharer/recipient pair, but this forfeits some bootstrapping benefits that FeedMe currently enjoys. We propose a third approach: build a receiver's public model by looking for terms that multiple sharers have all shared with that receiver. Only topics that are statistically "public knowledge" trigger such recommendations to a new sharer. We can also blacklist sensitive feeds or topics. Another concern is that social awareness load indicators, such as whether a receiver has read a post already, may leak sensitive content consumption information. We plan on giving receivers more control over what other sharers see.

CONCLUSION

Under threat of information overload, many people refuse to drink from the firehose of web content. These same people are interested in receiving more content than they do

now, however, and trust the information delivered by their social networks. As such, we seek to understand social link sharing so that we might enhance the process. We find that sharing is motivated by a perception of what friends would like to see, but held back by concerns about spamming and misreading friends' interests. We find that active information seekers are also the most active sharers, and have built a plug-in for such users of Google Reader to share over e-mail. Our plug-in, FeedMe, recommends friends who might be interested in seeing a post, reducing the effort required to share. The system highlights information relevant for sharers seeking to rate-limit themselves, and presents receivers with a lightweight thank-you mechanism. This work introduces a novel type of mixed-initiative social recommender system, where friends mediate recommendations rather than an artificial intelligence. FeedMe suggests that social sharing mechanisms offer a powerful new avenue for enhancing content recommendation on the web.

ACKNOWLEDGMENTS

We thank the participants in our study for their time and feedback. We also acknowledge Mark Ackerman, Katrina Panovich and our anonymous reviewers for suggesting many improvements to this work.

REFERENCES

1. Ackerman, M., Starr, B., and Pazzani, M. The do-i-care agent: Effective social discovery and filtering on the web. *Proc. RIAO '97*, (1997), 17–31.
2. Allen, T. *Managing the Flow of Technology*. MIT Press, Cambridge, MA, 1977.
3. Baumer, E. and Fisher, D. Smarter Blogroll: An Exploration of Social Topic Extraction for Manageable Blogrolls. *Proc. HICSS '08*, IEEE (2008), 155.
4. Baumer, E., Sueyoshi, M., and Tomlinson, B. Exploring the role of the reader in the activity of blogging. *Proc. CHI '08*, ACM Press (2008), 1111-1120.
5. Bernstein, M., Tan, D., Smith, G., et al. Collabio: A Game for Annotating People within Social Networks. *Proc. UIST '09*, ACM Press (2009), 177-180.
6. Burke, M., Marlow, C., and Lento, T. Feed me: motivating newcomer contribution in social network sites. *Proc. CHI '09*, ACM Press (2009), 945–954.
7. Dennis, B. and Jarrett, A. NusEye: Visualizing Network Structure to Support Navigation of Aggregated Content. *Proc. HICCS '05*, IEEE Press (2005), 107c.
8. Ehrlich, K. and Cash, D. Turning Information into Knowledge: Information Finding as a Collaborative Activity. *Proc. JCDL '94*, (1994).
9. Ellison, N.B., Steinfield, C., and Lampe, C. The benefits of Facebook "friends." *JCMC 12*, 4 (2007).
10. Erdelez, S. and Rioux, K. Sharing information encountered for others on the Web. *The New Review of Information Behaviour Research 1*, (2000), 219–233.
11. Erickson, T. and Kellogg, W.A. Social translucence: an approach to designing systems that support social processes. *TOCHI 7*, 1 (2000), 59-83.
12. Gamon, M., Basu, S., Belenko, D., et al. Blews: Using blogs to provide context for news articles. *Proc. ICWSM '08*, (2008), 60-67.
13. Granovetter, M. The Strength of Weak Ties. *American Journal of Sociology 78*, 6 (1973), 1360-1380.
14. Hearst, M. *Search User Interfaces*. Cambridge University Press, 2009.
15. Joachims, T., Freitag, D., and Mitchell, T. Webwatcher: A tour guide for the world wide web. *Proc. IJCAI*, AAAI (1997), 770–777.
16. Kittur, A., Chi, E., and Suh, B. Crowdsourcing user studies with Mechanical Turk. *Proc. CHI '08*, ACM Press (2008), 453-456.
17. Marshall, C. and Bly, S. Sharing encountered information: digital libraries get a social life. *Proc. JCDL '04*, (2004), 218–227.
18. Mason, W. and Watts, D. Financial Incentives and the "Performance of Crowds". *Proc. HCOMP '09*, (2009), 77 - 85.
19. McDonald, D. and Ackerman, M. Expertise recommender: a flexible recommendation system and architecture. *Proc. CSCW '00*, ACM Press (2000), 231–240.
20. Montaner, M., López, B., and La, J.L. A Taxonomy of Recommender Agents on the Internet. *Artificial Intelligence Review 19*, 4 (2003), 285-330.
21. Paepcke, A. Information Needs in Technical Work Settings and Their Implications for the Design of Computer Tools. *Computer Supported Cooperative Work 5*, 1 (1996), 63-92.
22. Resnick, P., Iacovou, N., Suchak, M., et al. GroupLens: An open architecture for collaborative filtering of netnews. *Proc. CSCW '94*, ACM Press (1994), 175-186.
23. Rocchio, J. Relevance feedback in information retrieval. In G. Salton, *The SMART retrieval system: experiments in automatic document processing*. Prentice Hall, Englewood Cliffs, NJ, 1971, 313-323.
24. Salton, G. and Buckley, C. Term-weighting approaches in automatic text retrieval. *Information Processing and Management 24*, 5 (1988), 513-523.
25. Sen, S., Geyer, W., Muller, M., et al. FeedMe: a collaborative alert filtering system. *Proc. CSCW '06*, ACM Press (2006), 89 - 98.
26. Terveen, L. and McDonald, D. Social matching: A framework and research agenda. *TOCHI 12*, 3 (2005), 401–434.
27. Terveen, L., Hill, W., Amento, B., et al. PHOAKS: A system for sharing recommendations. *CACM 40*, 3 (1997), 59-62.