Computer Science and Artificial Intelligence Laboratory

Technical Report

MIT-CSAIL-TR-2009-048 October 7, 2009

Understanding and Supporting Directed Content Sharing on the Web

Michael Bernstein, Adam Marcus, David Karger, and Rob Miller

massachusetts institute of technology, cambridge, ma 02139 usa — www.csail.mit.edu

Understanding and Supporting Directed Content Sharing on the Web

Michael Bernstein[1], Adam Marcus[1], David Karger, Rob Miller
MIT Computer Science and Artificial Intelligence Laboratory
32 Vassar Street, Cambridge MA 02139
{msbernst, marcua, karger, rcm}@csail.mit.edu

ABSTRACT

To find interesting, personally relevant web content, we often 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 present FeedMe, a plug-in for Google Reader that makes directed sharing of content a more salient part of the user experience.

Our survey research indicates that sharing is moderated by concern about relevancy to the recipient, a desire to send only novel content to the recipient, and the effort required to share. FeedMe allays these concerns by recommending friends who may be interested in seeing the content, providing information on what the recipient has seen and how many emails they have received recently, and giving recipients the opportunity to provide lightweight feedback when they appreciate shared content. FeedMe introduces a novel design space for mixed-initiative social recommenders: friends who know the user voluntarily vet the material on the user's behalf. We present a two week field experiment (N=60) demonstrating that FeedMe's recommendations and social awareness features 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.

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 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) readers 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 content that is interesting.

To improve this situation, we can build on the social interactions that help power information awareness today: we share web pages by e-mail, by talking in person, by posting to social networks, and more. 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 and tastes to forward you URLs. The social process has its own share of problems, however: friends can forward old or uninteresting material, and sharers worry about spamming their friends.

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 what drives social link sharing. We find that e-mail is still the dominant sharing medium, that topical interest in the content is the biggest determiner of recipient enjoyment, and that small numbers of recipients per link shared typically signal the most relevant content. We find that the best predictor of willingness to share content is an interest in seeking out new content, rather than a desire to increase social capital.

Armed with the knowledge that aggressive content consumers are also the most prolific sharers, we designed a tool to support such consumers in directed sharing of web content for the benefit those who do not want to drink directly from the firehose. To achieve this goal, we have developed FeedMe, a plug-in for the popular 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 to readers inline within RSS post viewing. Recommendations reduce the amount of effort required to share to two clicks: one click to select a recommended recipient, and one more to send. Social awareness and feedback help sharers avoid spamming by making visible information such as how much material the recipient has received today.

[1] Both students contributed equally to this work.

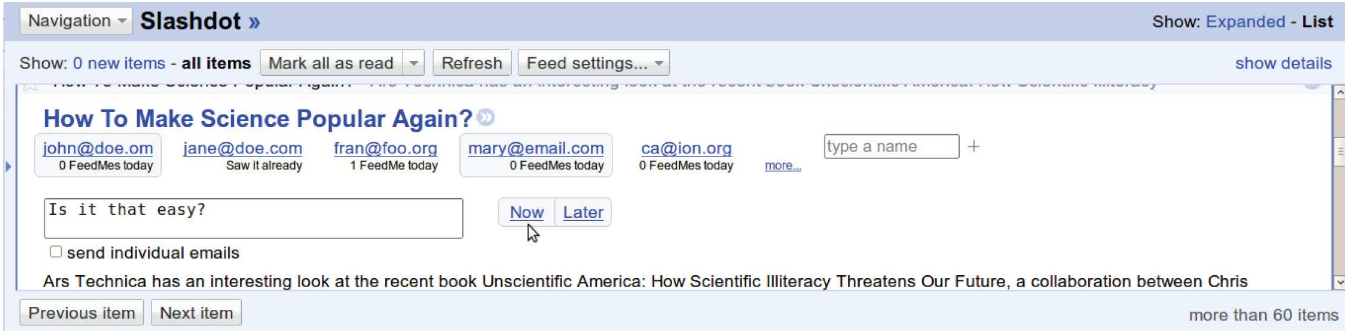

Figure 1. The FeedMe plug-in for Google Reader suggests friends, family, and colleagues who might be interested in seeing the post that you are reading. This user has selected john@doe.com and mary@email.com. The "Now" button sends an e-mail immediately; the "Later" button queues the item in a digest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

FeedMe also introduces a novel design space of mixed-initiative social systems. Most recommender systems suffer from two challenges: 1) they require training, and 2) they make mistakes. Our model is an extension of mixed-initiative interaction: a computer proposes, and a person acts on the recommendation. Instead of accepting the suggestions on their own behalf (like a search autosuggest helping to formulate queries), FeedMe allows the information seeker to act as a gatekeeper on another person's behalf. The recipient does not have to train FeedMe, since it rides on and facilitates a social process that is already happening. The system avoids mistakes because another person is vetting it.

In this paper we contribute: 1) an investigation of social link sharing practice as it stands: what motivates it now, and what keeps it from happening more; 2) a system that supports 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.

To follow, we review related work studying personalized information sharing, execute surveys of web users to study the social mechanisms behind sharing, present our plug-in FeedMe, and report on a two-week field experiment investigating FeedMe's impact on sharers and 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 mediator [8], Paepcke's contact broker [23] and Allen's technological gatekeeper [2]. 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]. Investigations of the microblogging platform Twitter and of collaborative search processes revealed that link-sharing is a central part of these practices [14, 22] as well.

We contribute to a growing set of literature studying and supporting blog reading. Baumer studied blog readers and found the practice to be a relaxing habit for respondents [4]. We contribute to knowledge of the role of the blog reader as a sharer, which, to our knowledge, has not been investigated despite the recent proliferation of sharing features on RSS readers.

FeedMe also takes inspiration from tools built to support blog-reading. Baumer and Fisher developed the Smarter Blogroll, which uses topic analysis to portray trending topics [3]. BLEWS is a visualization of discussion in the political blogosphere, identifying charged topics and liberal/conservative biases [11]. NusEye provides a graph layout for visual analysis of term co-occurrences in the blogosphere [7]. Other reader plug-ins such as PostRank (www.postrank.com) display posts of globally trending importance. FeedMe is the first tool to focus on directed sharing, and to do so with a mixed-initiative recommender approach.

One of FeedMe's core features is recommending recipients for interesting content, and we derive many of our web recommendation decisions from previous work in the field. Montaner et al. provide a taxonomy of Internet-based recommendation systems [21], and Schafer et al. characterize their role in e-commerce [27]. Popular techniques for recommendation interfaces include collaborative filtering (e.g., GroupLens [24]) and mining browsing history (e.g., WebWatcher [15]). The Do-I-Care Agent explored how collaboration technologies can support recommender systems, specifically by allowing agents to share information between themselves [1]. Rather than focusing on discovery, as do recommender systems, we focus on sharing; this necessitates the construction of profiles for users who are neither actively searching for content nor actively contributing to the content in their profile. Since FeedMe recommends *people*, we also build on work on expertise recommenders to match people to a piece of information (e.g., Expertise Recommender [20]). Unlike social matching [29], wherein individuals create their own profile and request to be matched with like-minded individuals, FeedMe builds on existing social connections, and requires no profile to start.

FeedMe uses humans to filter the content that a recipient accesses. One of the first projects to take this approach was that of PHOAKS by Terveen et al. [30], in which the system determines popular web content by crawling Usenet for frequent mentions of a given webpage. This approach has manifested itself more recently in services such as in Google News [6] and Digg (www.digg.com), but these approaches result in an un-personalized prioritization. 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.

Finally, we nod to other work which uses the FeedMe system name or paper title [5, 28].

SURVEYING EXISTING PRACTICE

Sharing interesting or amusing web content with others is woven into the fabric of web citizens' everyday social lives. To better understand the process, we extend previous research studying web sharing habits [10, 17]. First, we examine what *motivates* and *moderates* web sharing habits. What do recipients want to receive, and what do they consider spam? What concerns do sharers have when they consider forwarding a link? Second, we investigate the *characteristics of web content sharers*. Are sharers those who are social individuals, or instead those who voluntarily consume large amounts of web content?

We pursued these questions through a pair of online surveys. One survey serves as informal design research, painting a picture of the sharing process. The second survey quantitatively tests our hypotheses about the types of individuals who share the most web content.

How Does Sharing Happen?

We performed an initial survey to investigate how link-sharing happens today. 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 despised by recipients.

The *receiving* half of the survey investigated the process of receiving links from people whom one knows. We inquired whether receiving links 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 the broader social space of sharing, including whether forwards sent to groups had a different quality than those sent only to the individual, and which factors motivate a response to a link-share.

The *sharing* half of the survey investigated many of the same topics from the perspective of sharing rather than receiving. What are the significant motivators and fears when deciding whether to forward a link to a contact? What kinds of people are shared with most, and why?

We recruited 40 participants from Amazon Mechanical Turk to complete our survey. The survey consisted of a mix of multiple choice questions, free-response questions, and Likert scales. 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 methodology on best practices for Mechanical Turk user studies, such as making it equally difficult to answer our questions honestly and dishonestly via free-response [16]. We saw very little evidence of responders cheating the system; many wrote in-depth responses.

Results

Email is the Dominant Sharing Medium. For sharing and for receiving links we found e-mail to be the most common route. Ubiquity and consistency are important: everyone has an e-mail address and most people check their email constantly. When asked why they prefer e-mail as the method of receiving links, one participant reported "I am too busy for the other forms. I check e-mail throughout the day." Link sharing also is a face-to-face phenomenon: the topic would come up in conversation, and one person would show the link while both are present. These results support the conclusions of Erdelez and Rioux [10].

In addition, we investigated the how typical Internet users go about consuming information. Parallel to e-mail as a means of sharing, e-mail was also one of the most popular methods of encountering content. Directly visiting favorite web sites was another common pattern that almost all respondents followed. A smaller number of users utilized an RSS reader, or checked a social network for links.

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

Missing the mark was the most commonly cited reason for disliking interest- and entertainment-based URL sharing. Politics and YouTube videos were cited 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."

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 1. 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 to the person or group	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, and I do not want to overwhelm them	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 2. Being unsure of relevance to the recipient's interests is the largest concern cited with sharing. (N= 38)

Sharers are largely aware of their goals and the potential pitfalls. Of 39 respondents, 37 stated that their strongest motivator for sharing was the knowledge that the receiver would appreciate it (Table 1). In parallel, sharers' largest concern was determining whether the link would be relevant enough to the person or group (Table 2).

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 a particularly 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 2). 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, whose interests the sharers are relatively certain of. This is not a large group: sharers preferred sharing with individuals they are already in regular contact with.

Recipients Want Even 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." Recipients are willing to receive more messages: the median response was 6 out of 7 (μ=5.3, σ=1.3). This opens an opportunity for more sharing, if we can motivate the sharers to do so.

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 3. 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 4. Questions in the Sharing scale, investigating interest in passing web content on to others.

Who Are 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 as per our own experience and previous work [10, 17]: that *sharers are social individuals*, and that *sharers are people who spend a lot of time seeking out and experiencing web content*.

To operationalize these notions, we utilized existing scales measuring social capital and constructed two scales of our own: a *seeking* scale and a *sharing* scale (Table 3, Table 4). Social capital is a notion of how tied an individual is to the world around them, both to close friends and to the community. We adapted two social capital scales consisting of ten Likert-scale questions from Ellison et al. [9], measuring bridging social capital (how social the individual is in the community) and bonding social capital (whether the individual has a few strong ties [12] they trust). Our seeking scale measures how much time and interest an individual invests in finding interesting or funny 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.

We created a survey with all forty questions in random order for each participant. This survey was distributed via Mechanical Turk to a group of 100 individuals. Participants

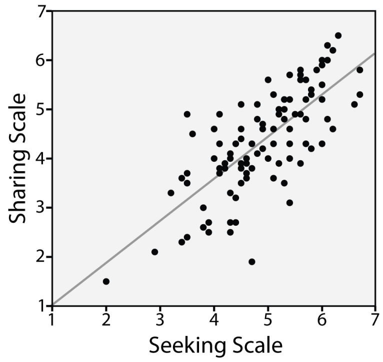

Figure 2. There is a clear relationship between an individual's expressed interest in seeking out web content, and their expressed interest in sharing web content. (N=100)

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 5. Interest in seeking has a very strong impact on interest in sharing, much more so than bridging social capital or bonding social capital. (Adj. R^2=0.56)

were again paid $0.20.

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 (Figure 2, Table 5). We found that interest in seeking is strongly correlated with interest in sharing (β=.74, $p < 0.001$), explaining more than half the variance in sharing scores. The social capital measures explain only 3% more variance when added to the model.

This survey indicates that an individual's interest in seeking out new web content is much more important than their sociality in determining how much they share web content with friends, family and colleagues.

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. We also cannot measure actual sharing counts.

Our sample may 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 demographics has given us reasonable confidence that at large N the results are reliable [19].

FEEDME

Our investigations suggested that we could increase the amount of productive, personalized sharing by pursuing two goals: 1) addressing concerns of whether a recipient might be interested in a post, reducing the effort barrier to sharing; 2) mitigating social concerns associated with the sharing process, like spamming and worrying whether a recipient already saw the link. We target our design at regular users of feed readers, as they have demonstrated interest in consuming lots of information and thus (as revealed in our surveys) are likely to be interested in sharing web content.

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: its recommendation interface, and its 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 select the name, 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 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.

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

Figure 3. FeedMe's load indicators change to reflect the number of FeedMe items sent today (left), or (right) when the recipient has received the post from another sharer.

a lightweight feedback mechanism.

Load Indicators

Our survey participants expressed fear about spamming contacts with too much content. To help the user gauge the likelihood of being perceived as spammy, FeedMe provides social awareness information in the recommendation interface (Figure 3). 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. Otherwise, the interface helps the sharer gauge how overwhelmed the recipient is by counting e-mails from FeedMe today. For example, if the recipient has received 2 FeedMe e-mails from one user and 3 from another, the interface displays "5 FeedMes today." We touch on the privacy implications of these indicators in the discussion section.

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 John Smith were to share a post, a link with the action text "Send John Smith 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, and is meant to foster friendly competition amongst receivers. 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 plugin for the Firefox web browser that facilitates the modification of a web site's code and interface. The script uses DOM listeners to determine when the user has shifted their attention to a new post. For each post, FeedMe sends AJAX requests to the FeedMe server for recommendations. The server is implemented using the Django web development 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 ever sent to the recipient, then tokenizes the result on spaces, performs word stemming, and removes common stop words. Words are weighted by term frequency-inverse document frequency (TF-IDF) [26], so that popular words in posts sent to the participant are more salient.

The recommendation algorithm uses the standard Rocchio approach [25], computing cosine distances to each friend of the sharer to the post and ranking the friends' distances. The FeedMe server generates a TF-IDF word vector per post, and ranks potential recipients by cosine distance to the post vector.

EVALUATION

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

We asked participants to fill 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 $t(158) = -3.215$, $p < .01$ respectively. 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 and publicly shared 5 posts. (Publicly sharing posts on Google Reader broadcasts a post to all of one's contacts who use Google Reader in a separate feed.)

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 (Figure 4). All factors were fully balanced.

Recommendations were either on, functioning as described, or as only an autocomplete textbox with no recommendations. This factor was within-subjects: participants tried each interface for a week.

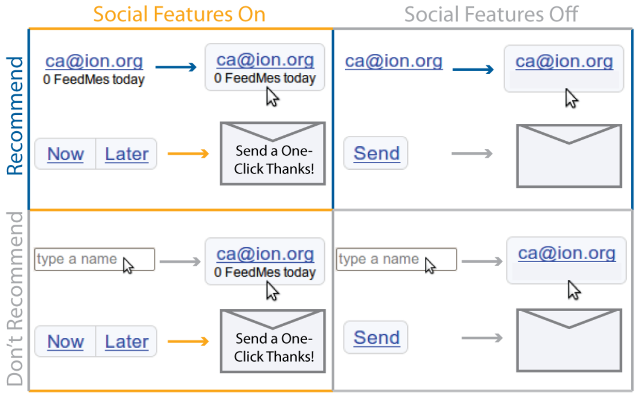

Figure 4. Feature combinations in our field experiment. Recommendations impacted whether the user needed to explicitly enter a name to share. Social features toggled the number of received items, the ability to digest with "Later," and the option for receivers to send One-Click Thanks.

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 because four (2x2) configurations would be more difficult for participants to remember and compare than two would.

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: 42% of our participants in the study continued to use FeedMe to share 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 sharers at ease, and digests freed some users to send many more posts than other study participants. To follow, we report these findings in detail.

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 84667 posts viewed while FeedMe was enabled in Google Reader. Figure 5 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, participants typically shared posts with a single recipient, and most participants shared 20 or fewer posts.

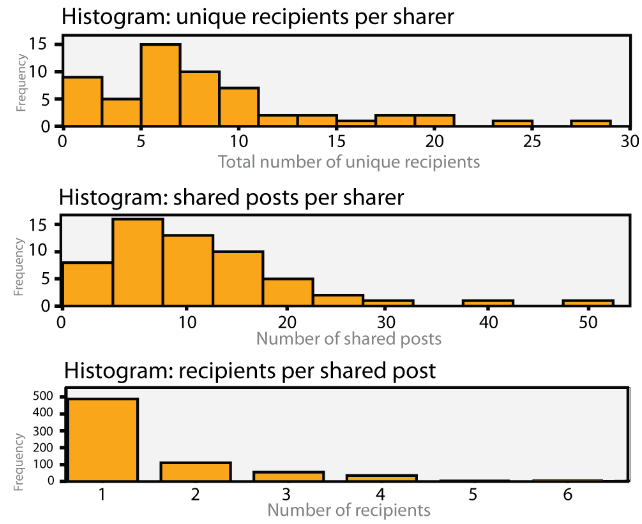

Figure 5. Typically, users shared with small numbers of individuals and addressed each message to one recipient.

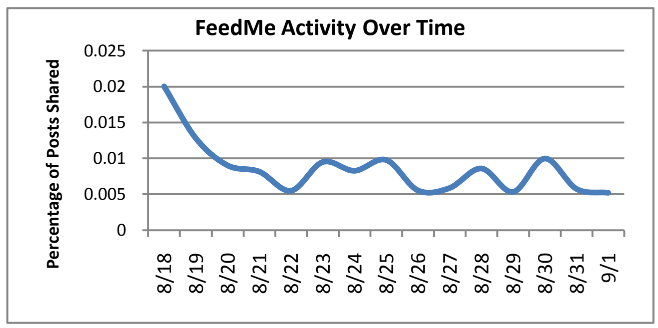

Figure 6. After the initial rush of activity, participants continued to use FeedMe to send a relatively consistent percentage of posts viewed throughout the two weeks.

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. After an expected initial flurry of activity, sharers shared a relatively constant number of posts per viewed article through the two weeks (Figure 6). 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.

All versions of FeedMe had a large effect on the amount of sharing occurring within the Google Reader interface. The median number of e-mailed posts from Google Reader in the 30 days prior to the study was 0. A paired t-test

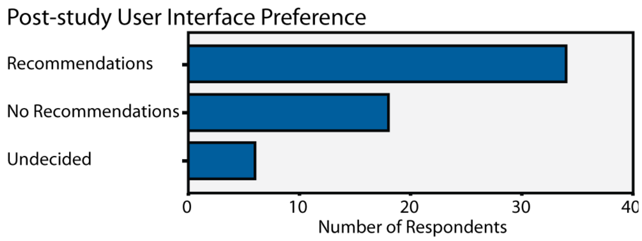

Figure 7. Participants reported a significant preference for the recommendation interface ($p < .05$).

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 within Google Reader.

Given the relatively small number of shared posts, we proceeded with our summative evaluation largely via qualitative and subjective assessments, augmented with usage statistics when relevant. 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 glad to have received the information: the mean Likert response was 5.1 (σ=1.6). Assuming a normal distribution of responses, this would indicate that for most posts, the receiver was neutral at worst and enthusiastic at best about the material.

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 received more information than they would have otherwise.

We conclude that recipients did not feel spammed by FeedMe, were pleased by the shared posts, and were better-informed 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. 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 [13], 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 7). 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 ($\chi^2(1, N{=}48) = 3.920$, $p < .05$), with nearly twice as many participants preferring recommendations to no recommendations (34 to 18).

Users who preferred the recommendation interface explained that recommendations made it easier to share articles. After each week's trial, we asked participants to rate on a Likert scale whether the version of FeedMe used made it easy to share items with recipients. To analyze noncontinuous Likert data, we utilized a nonparametric analogue to the paired t-test, the Wilcoxon Signed Ranks Test. The test indicated that recommendations made it easier to share ($z = 2.387$, $p < .05$). So, FeedMe's recommendations lowered the effort barrier to sharing.

Participants who preferred the no-recommendation interface did so for reasons of clutter and waste of vertical pixels in Google Reader. Some users expressed dissatisfaction with recommendations provided by FeedMe.

Acknowledging the small number of shares overall, we still wanted to investigate whether having recommendations actually led to more sharing. To test this hypothesis, we ran a repeated measures blocked ANOVA with recommendations (within-subjects) and social features (between-subjects) as independent variables, interface order as a blocking variable, and number of shared posts as the dependent variable. We dropped 10 participants from the statistics who did not log in to Google Reader at least once every two days as we asked. The effect for recommendations was not significant: $F(1, 56)=0.37$, *n.s*. It appears that the FeedMe recommendations do not strongly impact the amount of sharing that occurs. Our design efforts were targeted at the challenges identified in our survey research; having addressed the most significant ones, participants suggested that we faced a new barrier: concerns about inundating the un-ignorable e-mail medium. We postpone this discussion to the end of the study results.

Social Awareness and Feedback

Demand for the social features was high – participants who spent the two weeks without social features independently requested that they be added. Nine of the 30 users with social features mentioned digests, activity statistics, or One-Click Thanks as being their favorite feature in FeedMe. When asked what feature of FeedMe would make them feel more comfortable sharing more, 14 of 26 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 thirty users with social features made such a claim. The difference between these two groups is significant, as verified by a Mann-Whitney U between groups (Chi-Square(1, N = 48) = 10.08, $p < .001$). 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 an equivalent percentage (~30%) of posts on Facebook receive at least one Like—an equal engagement from a smaller audience. One FeedMe e-mail 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.

Returning to the statistical blocked ANOVA model test whether FeedMe's social features led to more sharing, we again found no significant main effect of social features on the number of posts shared ($F(1, 56)=0.57$, *n.s.*). One exception to this trend was the set of social-enabled users who chose to make use of the digest option. For some sharers, having access to a digest was liberating enough to lead to more sharing. An ANOVA with digest-user as the independent variable and number of shared posts as the dependent variable, blocked on participant and run only on active users in the social features condition, finds that digest users sent significantly more posts than users who did not make use of the digest option ($F(1, 20) = 4.40$, $p < .05$).

Opportunities for Improvement

The most vocal criticism of FeedMe 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 primary concern lay in the push mechanism underlying e-mail: recipients were forced to look at the links as a part of reading their e-mail. "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, which broadcasts to a feed for their Google Reader contacts to read at their convenience. 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. Rather than recommending content directly to a recipient, FeedMe recommends content to intermediate sharers who can efficiently and effectively verify that the receiver would be interested. Receivers never need to train the system, yet find the output of the combination to be enjoyable and novel. One benefit of this approach is that users are generally more tolerant of errors when acting on another person's behalf than when the recommendations are for themselves. We found that FeedMe's users appreciated recommendations when they were accurate, and generally did not mind if they were wrong. Another benefit of this design is the low marginal cost for sharers – they have already taken the time to find interesting content, and filtering requires relatively little additional overhead. We also note that the design space FeedMe opens up includes other domains where AI is still error-prone, for example expert finding: "We think that Sanjay can answer questions about nonparametric statistics—do you agree?"

FeedMe's field experiment highlighted issues with today's sharing media. We have presented several measures indicating that participants prefer a recommendation system for ease-of-use, and this system was designed directly from our survey research, yet there is no measurable impact on the number of posts shared with others. 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. Statistically, however, most people do not use IRC and RSS, and IM is viewed by many as another high-priority feed. Social network streams may soon provide an attractive alternative.

We designed FeedMe under the assumption that recipients would provide little feedback. We then found that recipients readily adopted lightweight social cues 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 relevance feedback from the recipients.

Privacy issues are worth addressing briefly. FeedMe's receiver models are accessible to all sharers in order to help bootstrap recommendations, but this has privacy implications if a friend is recommended an article on a sensitive topic (e.g., medical conditions) – another friend sharing posts with that recipient may stumble upon recommendations that would embarrass the recipient. To avoid this, we plan to build a receiver's public model out of the intersection rather than the union of sharers' models: only topics that are statistically "public knowledge" trigger such recommendations to a new sharer. We can also blacklist sensitive feeds, topics and terms, or maintain separate models per sharer. 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 sharers learn about them in the future.

CONCLUSION

Under threat of information overload, we rely on social solutions to identify web content that we will want to see. Our work seeks to understand the phenomenon of e-mail link sharing, the most popular means of sharing we surveyed, so that we might augment the process. We find that sharing is motivated by an understanding of what friends would like to see, but held back by concerns about sending undesired material. 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, lowering the effort barrier to share. FeedMe also highlights information relevant for sharers seeking to rate-limit themselves, and presents receivers with a lightweight thank-you mechanism. FeedMe represents a novel type of mixed-initiative social recommender system, where friends mediate recommendations rather than the end user. Our two-week field experiment reveals that FeedMe's recommendations and social features made sharing easier and more enjoyable. Receivers indicated that they were glad to receive the e-mails, and that most of the content was novel.

REFERENCES

1. Ackerman, M., Starr, B., Pazzani, M., et al. The do-i-care agent: Effective social discovery and filtering on the web. *Proc. RIAO '97*, 17–31, 1997.
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, 155-155, 2008.
4. Baumer, E.P., Sinclair, J., Hubin, D., et al. metaViz: Visualizing Computationally Identified Metaphors in Political Blogs. *Proc. SoSC '09*, IEEE, 2009.
5. Burke, M., Marlow, C., Lento, T., et al. Feed me: motivating newcomer contribution in social network sites. *Proc. CHI '09*, ACM Press, 945–954, 2009.
6. Das, A., Datar, M., Garg, A., et al. Google news personalization: scalable online collaborative filtering. *Proc. WWW '07*, ACM Press, 271–280, 2007.
7. Dennis, B. and Jarrett, A. NusEye: Visualizing Network Structure to Support Navigation of Aggregated Content. *Proc. HICCS '05*, IEEE Press, 107c-107c, 2005.
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., Lampe, C., et al. The benefits of Facebook "friends:" Social capital and college students' use of online social network sites. *Journal of Computer Mediated Communication 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. Gamon, M., Basu, S., Belenko, D., et al. Blews: Using blogs to provide context for news articles. *Proc. ICWSM '08*, 60-67, 2008.
12. Granovetter, M. The Strength of Weak Ties. *American Journal of Sociology 78*, 6 (1973), 1360-1380.
13. Hearst, M. *Search User Interfaces*. Cambridge University Press, 2009.
14. Java, A., Song, X., Finin, T., et al. Why We Twitter: Understanding Microblogging Usage and Communities. *Proc. WebKDD*, ACM Press, 56-65, 2007.
15. Joachims, T., Freitag, D., Mitchell, T., et al. Webwatcher: A tour guide for the world wide web. *Proc. IJCAI*, AAAI, 770–777, 1997.
16. Kittur, A., Chi, E., Suh, B., et al. Crowdsourcing user studies with Mechanical Turk. *Proc. CHI '08*, ACM Press, 453-456, 2008.
17. Marshall, C. and Bly, S. Sharing encountered information: digital libraries get a social life. *Proc. JCDL '04*, 218–227, 2004.
18. Mason, W. and Watts, D. Financial Incentives and the "Performance of Crowds". *Invited Talk An Overview of Human Computation*, 77, 2009.
19. Mason, W. and Watts, D. Financial Incentives and the "Performance of Crowds". *Proc. HCOMP '09*, 77, 2009.
20. McDonald, D. and Ackerman, M. Expertise recommender: a flexible recommendation system and architecture. *Proc. CSCW '00*, ACM New York, NY, USA, 231–240, 2000.
21. Montaner, M., López, B., La, J.L., et al. A Taxonomy of Recommender Agents on the Internet. *Artificial Intelligence Review 19*, 4 (2003), 285-330.
22. Morris, M.R. A survey of collaborative web search practices. *Proc. CHI '08*, ACM Press, 1657-1660, 2008.
23. Paepcke, A. Information Needs in Technical Work Settings and Their Implications for the Design of Computer Tools. *Computer Supported Cooperative Work 5*, (1996), 63-92.
24. Resnick, P., Iacovou, N., Suchak, M., et al. GroupLens: An open architecture for collaborative filtering of netnews. *Proc. CSCW '94*, ACM Press, 1994.
25. Rocchio, J. Relevance feedback in information retrieval. *The SMART retrieval system: experiments in automatic document processing*, (1971), 313-323.
26. Salton, G. and Buckley, C. Term-weighting approaches in automatic text retrieval. *24*, 5 (1988), 513-523.
27. Schafer, J., Konstan, J., Riedi, J., et al. Recommender systems in e-commerce. *Proc. eCommerce '99*, ACM Press, 158–166, 1999.
28. Sen, S., Geyer, W., Muller, M., et al. FeedMe: a collaborative alert filtering system. *Proc. CSCW '06*, ACM Press, 2006.
29. Terveen, L. and McDonald, D. Social matching: A framework and research agenda. *TOCHI 12*, 3 (2005), 401–434.
30. Terveen, L., Hill, W., Amento, B., et al. PHOAKS: A system for sharing recommendations. *CACM 40*, 3 (1997), 59-62.