

Talk Amongst Yourselves: Inviting Users To Participate In Online Conversations

F. Maxwell Harper¹, Dan Frankowski¹, Sara Drenner¹, Yuqing Ren²,
Sara Kiesler², Loren Terveen¹, Robert Kraut², John Riedl¹
CommunityLab*

¹University of Minnesota
200 Union Street SE
Minneapolis, MN 55455
{harper, dfrankow, mosch,
terveen, riedl}@cs.umn.edu

²Carnegie Mellon University
5000 Forbes Ave
Pittsburgh, PA 15221
{yren, kiesler, kraut}@cs.cmu.edu

ABSTRACT

Many small online communities would benefit from increased diversity or activity in their membership. Some communities run the risk of dying out due to lack of participation. Others struggle to achieve the critical mass necessary for diverse and engaging conversation. But what tools are available to these communities to increase participation? Our goal in this research was to spark contributions to the movielens.org discussion forum, where only 2% of the members write posts. We developed *personalized invitations*, messages designed to entice users to visit or contribute to the forum. In two field experiments, we ask (1) if personalized invitations increase activity in a discussion forum, (2) how the choice of algorithm for intelligently choosing content to emphasize in the invitation affects participation, and (3) how the suggestion made to the user affects their willingness to act. We find that invitations lead to increased participation, as measured by levels of reading and posting. More surprisingly, we find that invitations emphasizing the social nature of the discussion forum are effective, while invitations emphasizing non-social aspects of the discussion are less so.

ACM Classification: H.4.3 Communications Applications: Bulletin boards, H.5.3 Group and Organization Interfaces: Web-based interaction, J.4 Social and Behavioral Sciences: Psychology.

General terms: Design, Experimentation, Human Factors

Keywords: Online community, incentives, personalization, invitations, discussion forums

INTRODUCTION

Vibrant online communities offer ways for people with common interests to connect and organize their contributions for common purpose. Many online communities, especially those providing member-contributed content, contain both an archive of domain information and a social space where members exchange information and interact [18]. Typically, visitors first “lurk” at the periphery of the information space; later some of them become active participants in the social space [15].

Not all communities are equally successful. Some communities have a sufficient or excessive volume of posts, and may wish to encourage lurking over posting [15]. Other communities die from lack of participation. Butler found that over 50% of a large and diverse sample of email-based groups failed to receive a single message over the course of a four month study [4]. These communities might be helped by the presence of additional posters, who will contribute content, or even by additional lurkers, whose visible presence can help encourage contributions from other users [16]. MovieLens (<http://movielens.org>), the movie recommendation community which is the site of the research reported here, is one such community that would benefit from increased posting and lurking: only 19% of members visit its discussion forum, and only 2% write posts.

What tools are available to online communities to increase the level of participation? One class of tools is *designed incentives*, mechanisms built into a software interface that encourage or motivate users. It has become common for online systems to make visible some display of users’ reputation and grant enhanced system privileges to users who contribute more or better content. For example, Yahoo! Answers participants (<http://answers.yahoo.com>) earn “points” and “levels” for contributing answers of any quality, and Slashdot members (<http://slashdot.org>) are allowed

* CommunityLab is a collaborative project of the University of Minnesota, University of Michigan, and Carnegie Mellon University.
<http://www.communitylab.org/>

some powers as forum moderators and meta-moderators after earning enough “karma” from activities such as writing high quality posts.

In this research, we explore *personalized invitations*, a particular type of designed incentive. Personalized invitations encourage members to visit or post in a discussion forum by augmenting the user interface to emphasize the presence of interesting content. The success of personalized invitations relies on intelligent computation to generate appropriate content for display, as well as appropriate presentation to maximize the visibility and potential effectiveness of the appeal.

There is reason to believe that displaying invitations to users may lead to action. A fundamental principle in human behavior is that people do things to minimize their behavioral costs or effort. Zipf [23] identified this principle in the 1940s, using it to account for humans’ tendencies to develop shorter words as the words become more frequent in a language (e.g., “television” compressing to “TV”), to communicate most with people who are close by, and to select the pie closest to the front of the freezer. This principle helps explain why people make decisions heuristically, rather than through a rational analysis of costs and benefits [20], and why they use heuristic processing of persuasive messages rather than the more systematic analysis of the evidence that a message presents [5].

We further believe that the use of personalization to tailor the content of an invitation to a particular user will improve on a non-personalized call to action. Prior work in recommender systems has shown that personalization can help users make decisions when faced with uncertainty [10]. E-commerce Web sites have used this knowledge to build personalized interfaces to increase sales [19]. In addition, Internet users claim to prefer personalized content to non-personalized content [9].

We measure the success of personalized invitations by the degree to which they increase the amount of activity in a discussion forum. Since this research is focused on communities that suffer from too low a level of activity, the amount of activity in the forum is an important indicator of success. We measure total activity in terms of posts read and written.

In the following sections, we describe and evaluate algorithms and interfaces for delivering invitations. At the highest level, we find:

Invitations emphasizing the social nature of the discussion forum increased user activity, while invitations emphasizing other details of the forum were less successful.

The results from this research can be used by designers who wish to increase participation in a discussion forum. Intelligent algorithms can relate users in the community with one another to as a way of breaking the ice and encouraging new relationships. And new interfaces can emphasize social

features such as the names of other users and recent post text to emphasize the presence of interesting people and discussions.

RESEARCH QUESTIONS

We have three questions we wish to answer:

RESEARCH QUESTION 1: OVERALL. *Do personalized invitations lead to increased participation?*

RESEARCH QUESTION 2: ALGORITHMS. *Do different algorithms for choosing the content of the invitation affect users’ response rates?*

RESEARCH QUESTION 3: SUGGESTIONS. *How does the suggestion made in the invitation affect users’ willingness to act?*

These questions are broad, and could be investigated in a variety of ways. We have chosen to pursue our investigation in the context of MovieLens, because this site offers us the opportunity to run controlled field experiments with large numbers of users. MovieLens offers an abundance of data about users’ movie preferences, and allows us to track both posting and reading behavior.

In the subsequent section, we describe MovieLens and its discussion forum, in order to provide some sense for the opportunities for designing personalized invitations. We follow this with a description of the general framework for generating personalized invitations. Finally, we report on two iterative field experiments in which we evaluate several variants of personalized invitations that represent interesting points in the design space.

RESEARCH CONTEXT

MovieLens is a movie recommendation web site where people rate movies and receive movie recommendations. The site has more than 100,000 registered users, and averages over 2,000 unique visitors each month. MovieLens collects movie ratings from its users, and uses automated collaborative filtering [17] to generate personalized recommendations.

In June 2005, MovieLens was augmented to include a discussion forum. MovieLens’s forum (based on open-source mvnForum) is a non-hierarchical, threaded conversation space with two main areas of conversation: one for talking about movies, and the other for talking about MovieLens. These forums are publicly visible, but posting requires (free) registration.

While the discussion forum has attracted some dedicated users, the number of regular forum users is a small fraction of the total number of MovieLens users. Of the approximately 12,000 unique members who have visited MovieLens since the launch of the forums, only 19% have visited the forums, and only 2% have posted one or more messages. Of the people who read the forums, 88% are lurkers (people who read but do not post). Lurkers may help increase participation, by motivating posters [16]. MovieLens

shows the presence of lurkers by displaying read counts next to threads, and by showing a list of online users on the front page of the forum.

One distinguishing feature of the MovieLens forums is its ability to recognize and understand references to movie titles in posts [7]. Recognized movies are hyperlinked and recommendation information is presented alongside the forums interface. Figure 1 shows a post with two linked references. Movie recognition enables several of the personalization algorithms described in this research below.



Figure 1. Forum post with two movie references.

Movie recognition is enabled by two custom tools, the *movie-linker* and the *movie-finder*. The movie-linker allows members to manually insert a movie reference into the text of their post, using an AJAX-based completion interface for quick movie searches in-place. The movie-finder automatically inserts linked references when titles are found in the content of the post. Our approach to automatically discovering movie entities in post text is related to prior work done in the field of natural language processing (e.g. [14]), although our focus has been on improving the usability of named entity recognition systems. The movie recognition interface and architecture used in MovieLens are described in more detail in [7].

THE DESIGN SPACE

Personalized invitations are persuasive messages that encourage members to visit or post in a discussion forum. While the potential design space for invitations is large, we consider two main aspects: the algorithms used to select content for display, and the nature of the suggestion made to the user. We do not investigate other potentially interesting aspects, such as visual design or the timing of the presentation.

Figure 2 gives a conceptual overview of the design space we consider. An invitation, shown in the lower-right corner, has two components. *Content* refers to a system entity (e.g., a forum post, a movie, or a user name) that comprises the focus of the invitation. *Suggestion* refers to the request made of the user, and the presentation of the request. While a suggestion is coded into the invitation’s design, content is dynamically generated. To find content, a *content selection algorithm* searches through system data for entities to display, possibly based on the target user’s preferences or usage history.




Figure 2. Conceptual Diagram of the Problem Space.

Figure 3 shows an example of a personalized invitation on the MovieLens home page. The content selection algorithm (*Rare Rated*, see below) has found a movie entity (*Nicholas Nickleby (2002)*) that the target user has rated, but that few other users have rated. The invitation suggests that the user start a new thread about this movie, stating that the user is “one of only a few” members who is able to take this action.



Figure 3. MovieLens home page with one of the invitation variants from experiment 1 (magnified).

In the following section, we describe the first of two iterative experiments on the effect of personalized invitations. We propose and evaluate several algorithms for intelligently choosing content for display, and examine the effectiveness of various suggestions for action.

EXPERIMENT 1: INVITATIONS TO POST

Our first experiment tested the overall effectiveness of invitations that ask users to post messages in the MovieLens discussion forum. Invitations designed for this experiment asked a user to either start a new thread about a specific movie or to reply to an existing post. We developed several algorithms for recommending posts or movies to supply the topic of the invitation.

People universally seem curious to learn about themselves and how they are similar to and different from others [8], suggesting invitations that emphasize ways that a person

differs from other members of the community. Prior work by Ludford et al. showed that participants in a discussion forum were motivated by knowing what unique perspective they can bring to a group, and that knowledge increases their participation [13]. Thus, a target user’s uniqueness is something that designers might choose to emphasize in an invitation. For example, in MovieLens, we might point out that a user can write about a movie that few other members have seen, or that a user’s opinion about a movie is unique. We extend the work of Ludford et al. by building algorithms that automate the selection of content that emphasizes uniqueness, and by evaluating the algorithm’s potential for enhancing the effectiveness of invitations to participate.

Extrapolating from the law of least effort [23], we can expect to elicit more contributions if we minimize the effort to participate. One way to vary the effort involved in writing a post is to either ask users to reply to an existing post, or to ask them to start a new topic of conversation. We hypothesize that replying to a post is a lower cost action than starting a new thread. By offering a specific post to reply to, we constrain the space of things to talk about. We test this idea by comparing the effectiveness of invitations asking users to start a new thread with invitations asking users to reply to an existing post.

Content Selection Algorithms

We recommended two types of content that have the potential to make an invitation compelling in MovieLens – movies and forum posts. The trick is choosing the right movie or the right post.

Recommending Movies. The MovieLens movie recommendation system offers many plausible ways for choosing movies to include in an invitation, such as: recently-rated movies, highly-rated movies, poorly-rated movies, rarely-rated movies, highly-recommended movies, and movies that are influential to a user’s recommendation model. Cosley et al. investigated several of these algorithms to find movies for users to edit for accuracy in MovieLens. They found that the choice of algorithm had a large effect on users’ willingness to help [6].

Based on the best algorithm from Cosley et al.’s investigation [6], we developed *Rare Rated*, an algorithm designed to choose a movie the user will be able to write about in the forums. This algorithm searches the target user’s ratings history for rarely-rated movies – defined as movies with fewer than 250 system-wide ratings – and picks one to display at random. As a baseline algorithm, we also developed *Rated*, an algorithm that chooses a movie for display at random from the target user’s rating history.

Recommending Posts. There are a variety of ways to choose forum posts for inclusion in an invitation. Surprisingly, there is little in the research literature that discusses the problem of recommending posts that users may be interested in reading. But one can imagine many plausible algorithms for this task, ranging from collaborative filtering

to content-based algorithms.

One approach to choosing forum posts for display relies on what we call *indirect recommendation*: the use of knowledge in one domain to recommend items in another domain. In MovieLens, we can choose posts for display in an invitation by combining our knowledge of users’ movie tastes with our knowledge of which movies have been mentioned in the forum.

There are several ways to use indirect recommendation to choose posts for display. We may recommend posts that mention movies a user is *familiar* with, that mention movies a user *likes* or *dislikes*, that mention movies the poster and the recipient of the invitation *agree* or *disagree* about, and that are written by authors who are *similar* or *dissimilar* to the recipient of the invitation.

For this experiment, we developed *Disagree*, an algorithm which recommends forum posts for users to reply to. This algorithm searches for posts referring to movies the target user has rated. Each post is scored on the maximum difference between the user’s rating and the post author’s rating of a movie referenced in the post. The system chooses the post with the largest disagreement score for display in the invitation. We developed the baseline algorithm *Random Rated Post*, which randomly chooses among all posts which refer to a movie the target user has rated.

Invitation Variants

We designed four variants of a personalized invitation in this experiment. These variants shared the same basic visual design: each contained a link to a movie the user had rated and a link to a page for writing posts. Figure 3 (above) shows a screenshot of the MovieLens home page, with an example invitation highlighted. The four invitation types were:

1. *New Thread + Random Rated.* This invitation asks MovieLens users to start a new thread about a movie. The system randomly chooses a movie the user has rated for display in the invitation. This invitation is worded as follows:

“Tell Others About [MovieTitle] – **You have rated the movie** [MovieTitle]. Post and share your thoughts.”

2. *New Thread + Rare Rated.* This invitation asks MovieLens users to start a new thread about a movie, emphasizing their ability to make a unique contribution by selecting content with the *Rare Rated* algorithm. We modify the bolded portion of the previous wording to state “you’re one of only a few MovieLens members who have rated...”.

3. *Reply + Random Rated Post.* This invitation asks users to reply to a post in the forums. The system searches for posts that refer to movies the user has rated, and chooses one at random. This invitation is worded as follows:

“Tell Others About [MovieTitle] – A recent post mentions the movie [MovieTitle], which you rated [Your-

Rating]. Post and share your thoughts.”

4. *Reply + Disagree*. This invitation asks users to reply to a post in the forums, emphasizing the potential of the user to provide a new perspective to the discussion by selecting content with the *Disagree* algorithm. We modified the wording above to include the sentence “We think you disagree with the poster about this movie.”

Methods

Our subjects were drawn from the pool of new and returning MovieLens members during 17 days in December 2005. We randomly chose 1/5 of the MovieLens members as the control group; the remaining 4/5 were assigned to the experimental group.

The control group continued to use an unmodified MovieLens interface for the duration of the study. The unmodified interface includes several paths to the forums: every page contains a header with a link to the forums, the front page contains a list of links to three recent forum posts, recently mentioned movies in search results link to forum posts, and each page devoted to a movie links to recent relevant forum posts.

The experimental group used the same MovieLens interface as the control group, with the addition of an invitation at the top of three pages: the MovieLens home page, the forums front page, and the page listing the threads in each forum. Each time a user in the experimental groups viewed one of these pages, one of the four invitation types was chosen at random, personalized, and displayed.

Results

The experimental group consisted of 1,611 users who logged in a total of 6,392 times. The control group consisted of 410 users who logged in a total of 1,552 times. 19% of the experimental group viewed at least one post, as compared with 20% of the control group. 3% of the experimental group posted at least once, as compared with 1% of the control group. 9% of the experimental group clicked on an invitation at least once.

The MovieLens pages that contained invitations in the experimental condition were viewed 27% more often by subjects in the experimental condition than by subjects in the control condition. Recall that in this study, each time a user views a page with an invitation, the content of the invitation is changed. We hypothesize that users viewed these pages more often because they were intrigued by the presence of the invitations, and acted to explore the different invitation types.

Overall Participation. Table 1 shows the mean number of posts written and viewed per user. To test for significance we used the non-parametric Kruskal-Wallis test. Users in the experimental group viewed (means = 15.2 vs. 11.8, $p < .01$) and wrote (means = .17 vs. .04, $p < .01$) posts more often than users in the control group.

	Control Group	Experimental Group
Posts Written	0.039	0.168
Posts Viewed	11.868	15.240

Table 1. Mean number of posts written and viewed per user.

Finding 1. Invitations increased participation.

New Thread vs. Reply. Table 2 summarizes the total number of invitation views, clicks, and posts over the duration of the experiment, across all users.

	Views	Clicks	Posts
New Thread	2397	17	9
New Thread + Rare Rated	2500	48	5
New Thread Subtotal	4897	65	14
Reply	2480	65	1
Reply + Disagree	2446	172	5
Reply Subtotal	4926	237	6

Table 2. Total views, clicks, and posts by invitation type. Only posts directly caused by an invitation are counted as posts. It is impossible to know exactly how many posts were indirectly caused by the presence of the invitations.

Reply invitations were clicked 237 times, while *New Thread* invitations were clicked 65 times. This difference is statistically significant, based on a logistic regression model built to predict whether or not a user clicked on the invitation ($z=4.84$, $p<.01$). However, *New Thread* invitations directly led to 14 posts, while *Reply* invitations only led to 6 posts. There was not enough data to detect a statistical difference.

Finding 2. Users were more likely to click on invitations asking them to reply to a post than on invitations asking them to start a thread.

Table 3 summarizes the number of clicks per user, grouping invitations by algorithm type. Across both the *New Thread* and *Reply* invitation types, users clicked on invitations with content emphasizing their uniqueness (*Rare Rated* and *Disagree*) more frequently than on invitations using the baseline algorithms (*Random Rated* and *Random Rated Post*). The effect of these algorithms on users’ click rates is positive and statistically significant, tested with a logistic regression model built to predict whether or not a user will click on an invitation ($z = 3.52$, $p < .01$).

	Low Uniqueness	High Uniqueness
New Thread	.0036	.0099
Reply	.0137	.0401

Table 3. Invitation clicks per user by invitation type. High Uniqueness counts the *Rare Rated* and *Disagree* algorithms together, while Low Uniqueness counts the two baseline algorithms together.

There was no difference between the high uniqueness and

the low uniqueness algorithms in terms of posting behavior.

Finding 3. Invitations emphasizing uniqueness led to more clicks, but not more posts.

Discussion

Research Question 1: Overall. Compared with the control group, users who received invitations read more and posted more, as shown in Finding 1. It also appears that invitations increased the overall activity level in the forum. Although the fluctuations of posting levels contain too much noise to draw statistical inference, there were on average more posts per day during the experimental period (19.5) than during the two weeks immediately before (13.7) and after (12.5) the experiment.

Did invitations spark valuable content, or did they lead to lower quality posts? Although quality is subjective and difficult to measure, we propose that one sign of a post's quality is whether it leads to further conversation. Of the 293 posts written during this experiment, the 20 posts directly caused by invitations have been replied to an average of 1.7 times, as compared with an average of 0.8 replies for the remaining 273. 75% of the posts directly caused by invitations received at least one reply, as compared with 64% of the other posts. Thus, invitations led to posts that were read and replied to.

However, there were also undesirable outcomes. In one instance, a new post author created a thread about a movie that had already been discussed elsewhere – an action which more senior members noticed and corrected. Lampe and Johnston suggest that old and new members alike benefit when new users spend time learning the customs of the community by reading posts before they first post themselves [11]. Thus, in our second experiment, we adjusted our suggestions to ask users to read posts, in order to give new users a chance to acclimate to the community before being thrust into the position of posting.

Research Question 2: Algorithms. Finding 2 shows that the algorithm for choosing content does matter. For both invitation types, personalization that emphasized uniqueness made a positive difference to response rates. One user commented on the personalization:

It's kind of interesting to see what movies pop up... and that I am "only one of a few" to have rated the movie.

The *Reply + Disagree* invitation variant was especially effective at generating clicks. It led to nearly three times as many clicks as the *Reply + Random Rated Post* variant.

Research Question 3: Suggestions. Finding 1 presents conflicting results concerning fundamental issues of presentation. While users were more likely to click on an invitation asking them to reply to a post, they were more likely to write a post when asked to start a new thread. Although the post-writing result was not significant, it remains a surprise. Replying to a message intuitively feels like a lower-cost action than writing a new thread, because the topic of con-

versation is more constrained. Based upon feedback from MovieLens users, we believe that this data may be explained by a shortcoming in the algorithms used to find posts for the *Reply* invitation variants: we asked users to reply to potentially very old posts (up to six months old). One user wrote:

Don't ask me to reply to a post that's more than a month old. (Better still, two weeks.)

In the next experiment, we address this concern by filtering recommended content for recency.

Next Steps. One of the most significant effects from experiment 1 was the increased activity generated by the *Disagree* algorithm. This invitation variant led to nearly three times as many clicks as the baseline invitation asking users to reply to a post that mentioned a movie. It may have been successful by creating a sense of curiosity about other users: “who do I disagree with about this movie, and why do our opinions differ?” We hypothesize that the *Disagree* algorithm made more visible the social nature of the discussion forum. Therefore, in experiment 2, we explore other invitations which emphasize the social nature of the forum.

EXPERIMENT 2: INVITATIONS TO READ

We revisited our design of personalized invitations in our second experiment, based on the experimental findings from the first experiment and feedback from users. We wished to study other aspects of invitation design. First, we wished to design invitations that bring users into the forums as readers, to give new users the chance to explore before they post. Second, we wished to understand other social dimensions of algorithms for choosing content.

People use familiarity to reduce cognitive effort. Habit enables people to make repeated decisions without having to think through the alternatives each time. People reduce the possibility of incorrect or suboptimal decisions by seeking familiar sources. Studies show that “mere exposure” [22] can explain people's attraction not just to other people, but also to music, art, and food. Repeated contact over time causes a person to like other people and objects more. These observations suggest that invitations presenting familiar items or people will be more effective than invitations that present unfamiliar items or people. Thus, in this experiment, we either chose familiar or unfamiliar content for display in the invitation.

How much should we reveal about the content of the invitation? The invitation might be more or less specific about the recommended content. For example, we might choose to display an entire post, or we might simply show the subject line. The study of the psychology of curiosity (e.g. [12]) posits that motives often stem from incongruities or information gaps in the world. If we display less information about a post, will that enhance users' curiosity? Or, will displaying less information cause us to omit details that are especially compelling to the user, such as a the title of a favorite movie, or an especially intriguing username. We test these ideas by varying the specificity of the information

displayed in the invitation.

Content Selection Algorithms

Because we wished to emphasize social aspects of the discussion forum, we exclusively choose post content for inclusion in invitations. We especially were interested in examining the power of emphasizing the social nature of the forums. Thus, we developed two algorithms, one that is intended to emphasize the social, and one that is intended to emphasize the unsocial.

Our social approach to choosing posts relies on users' history of viewing the discussion forum. By tracking which users have viewed which conversation threads, we can compute a familiarity score from any user to any other user. The *Relaviz* system [21] used a graphical display of the asymmetrical familiarity between pairs of users to encourage participation by connecting lurkers and posters. We use a similar computation to discover relationships between users, but use the output to recommend posts written by *familiar* or *unfamiliar* users. The recommendation of familiar users is only possible for MovieLens members who have previously viewed forum posts.

We call our algorithm that implements this idea *Familiar Poster*. This algorithm chooses among posts written in the last week based on whether or not the target user has previously viewed posts by that author five or more times. Requiring fewer views weakens users' familiarity with post authors, while requiring more views further restricts the set of posts to recommend; we looked to balance these two factors. We developed a corresponding baseline algorithm, *Unfamiliar Poster*, to recommend posts written by authors the target user has seen fewer than five times.

To test whether any effects of familiarity were due to the social effects of making other users visible, we also developed a non-social approach to recommending posts with familiar or unfamiliar content. Just as we may believe that users are attracted to the forum by the presence of other users, we may also believe that users are attracted to the forum by the movies that are discussed.

Our algorithm for recommending posts on the basis of their movie content is called *Familiar Movie*. It returns the set of posts written in the last week that reference one or more movies the target user has rated. We developed a corresponding algorithm, *Unfamiliar Movie*, that returns the set of posts which mention movies the target user has not rated. While a single post may contain both rated and unrated movies, in the invitation we only show the movie chosen by the algorithm.

Invitation Variants

We made several changes to the overall design of invitations in MovieLens in preparation for our second experiment. The biggest change was that the invitations no longer asked users to post, but instead recommended that users visit the forums to read a post.

Some subjects thought that getting many different invitations per session (the time between login and logout) was too demanding in the first experiment. Thus, we designed these invitations so that users would view the same invitation throughout their session. Clicking removed the invitation from the user's interface for the duration of the session. Users could also explicitly hide the invitations for the remainder of the session using a "hide" link. Users who had clicked "hide" in three or more sessions were given a "hide forever" link that permanently removed invitations from their view when clicked. 65 out of 1,917 users clicked the "hide" link at least once; 11 users chose to permanently hide invitations, of whom 6 are repeat forum posters.

Figure 4 shows a sample invitation. The most prominent visual change to the invitation design was the inclusion of the subject line of the recommended post, as well as up to 125 characters of preview post text. To avoid confounding our experimental manipulation, we stripped references to movies from the preview text. To do this in a natural way, we searched the post for the first phrase beginning with 125 characters without a movie reference. Failing this, we used the first 125 characters of the post, replacing movie titles with the string "...".



Figure 4. A sample invitation from the MovieLens home page during experiment 2. The post preview text is used with the author's permission.

There were a number of variants of this basic invitation, as described below.

Specificity Variations. We call invitations that contain more information about the recommended post more *specific*. Users were either shown the name of the author of the recommended post, or they were not. Similarly, users were either shown the name of a movie referenced in the recommended post, or they were not. We randomized the order of the name and the movie in invitations that displayed both. Named entities were shown in bold, green text to draw attention to their presence. We added small icons next to these named entities to further distinguish their presence. We adjusted the two icons to be the same size and approximately the same level of luminosity.

Familiarity Variations. We also varied whether or not the invitation's named entities were *familiar* to the user. In the case of a movie, we varied whether we used the *Familiar Movie* or the *Unfamiliar Movie* algorithm to select the content. In the case of a post author, we either used the *Familiar Poster* or the *Unfamiliar Poster* algorithm. We randomly chose among the intersection of these sets of posts for display.

Credibility Variations. We designed four different wordings to use as the opening phrase in the invitation. Two of the wordings were designed to exploit the *credibility* of the MovieLens recommender system by implying that “our system” is making a recommendation for the user.

Our system predicts you'll enjoy the following new post

Our system recommends the following new post for you

The other two wordings did not imply that it is the system making the recommendation, but left the source of the recommendation unspecified:

We think you'll like the following new post

Check out the following new post

Methods

The second study took place for 17 days in February and March, 2006. All MovieLens users who logged in during this period and who met our entrance criteria were exposed to the experimental manipulation. We did not include a control group because we had established the efficacy of the invitations in the first study, and because we were primarily interested in measuring the effect of experimental manipulations.

For the purposes of analysis, we split user sessions in this experiment into two groups:

1. *ForumHistory*: users who have visited the forums in a previous session
2. *NoForumHistory*: users who have never visited the forums

Users in *ForumHistory* received the full set of invitation variations, while users in *NoForumHistory* received all variations except for those that require a familiar post author. Users were moved from group *NoForumHistory* to group *ForumHistory* after their first visit to the discussion forum. Users who could not receive the full set of invitation variations did not receive any invitation, and these sessions are not included in the analysis. For example, users who had not rated any of the movies that had been referenced in the past week of forum posts could not receive the familiar movie variation, and were excluded.

Due to the exploratory nature of our study and a limited pool of users, we used a half fractional factorial design, a design that gives us main effects and lower-order interactions, but sacrifices higher-order interactions. We chose 8 out of 16 runs of a full, four-factor, two-level, factorial design [2]. The factors are: showing a movie title, familiarity with the movie title, showing a user name, and familiarity with the user name. The levels are: *ForumHistory* and *NoForumHistory*.

Our data are nested by nature, since each user can have multiple sessions. We analyzed the data using hierarchical

linear modeling (HLM) techniques [3] to control for random effects at the user level. We then determined the significance of our results using HLM analyses.

Results

2,415 users logged in to MovieLens during experiment 2. 1,917 of these users participated in the experiment by viewing an invitation. 10.5% of the participating users clicked on an invitation at least once. Table 4 summarizes the number of users, sessions, and invitation clicks that took place during the study.

	Unique Users	Sessions	Clicks
ForumHistory	704	3012	193
NoForumHistory	1213	3225	80
Total	1917	6237	273

Table 4. Number of users, sessions, and invitation clicks in experiment 2 by user group.

Overall, users in *ForumHistory* clicked on an invitation in 6.0% of their sessions, as compared with users in *NoForumHistory*, who clicked in 2.5% of their sessions. This effect is significant in our model ($p < .01$). This effect might be expected, given that users in this group had visited the forums in the past.

Credibility. Table 5 summarizes the percentage of invitations that were clicked by wording. The two wordings that contained the phrase “our system” led to users clicking on the invitation 50% more often than with the other two, a significant effect in the regression model ($p < .01$). Overall, the most effective wording was “Our system predicts you'll enjoy the following new post,” with 99 clicks. The least effective wording was “Check out the following new post,” with 51 clicks.

Wording	% Clicked
Our system predicts you'll enjoy the following new post	5.4
Our system recommends the following new post for you	5.0
We think you'll like the following new post	4.1
Check out the following new post	2.6

Table 5. Percentage of invitations clicked for each wording, across all groups. The first two wordings emphasize the credibility of the source of the recommended post.

Finding 4. The wording of the invitations mattered. Invitations were more effective when they were worded to emphasize the credibility of the recommendation.

Specificity: Usernames. Simply showing the name of a forum poster in the invitation has a positive and significant statistical effect on click rates ($p < .01$). Importantly, though, this effect only applies to users in *ForumHistory*, as shown in Figure 5. The interaction effect between user group and showing the poster’s name is statistically significant in our model ($p < .01$).

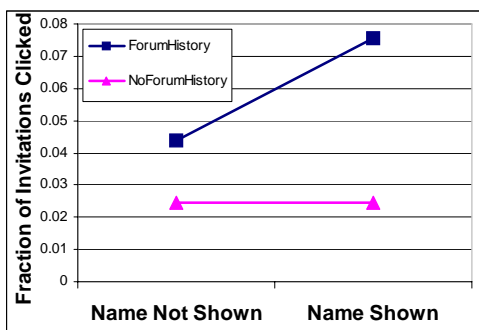


Figure 5. Effect of poster’s name on click-through rate by user group.

Finding 5. Invitations were more effective when they showed the name of a post author, but only for users who had previously visited the forum.

Specificity: Movies. Showing the title of a movie in the invitation had ambiguous effects. It actually slightly depressed the click rate for users in *NoForumHistory* (from 3.3% to 3.1%), although it increased the click rate for users in *ForumHistory* (from 5.6% to 6.5%). Neither change is statistically significant.

Finding 6. Invitations were not improved or worsened by displaying the name of a movie.

Familiarity. We failed to find evidence that our personalization algorithms for detecting familiar movies and forum authors improved the success of our invitations. While we initially saw evidence that suggested showing familiar movie titles increased click rates (from 4.0% to 5.4% across groups), further analysis showed that our findings were confounded. Whether or not a user was familiar with a movie is correlated with the total number of times the movie has been rated in MovieLens (correlation = .448). Thus, it is impossible to know if the increased click rate is due to the user’s familiarity with the movie or its overall popularity. When both factors are included in our model, neither is significant.

Likewise, we do not find evidence that showing familiar user names affects click rates. While it is the case that click numbers were slightly higher when a familiar name was shown (8.3% click rate when the name was familiar vs. 6.9% when the name was unfamiliar), the effect is not statistically significant.

Finding 7. Invitations were not improved by familiarity-based personalization.

Discussion

Research Question 2: Algorithms. Our algorithms for choosing posts on the basis of entity familiarity (or lack of familiarity) had no effect on users’ rates of clicking through to the discussion forum. Perhaps this finding (Finding 7) points to a fundamental tension between familiarity and curiosity: while familiar people and items may be more comfortable and liked, unfamiliar people and items may

heighten a target user’s curiosity. For example, we might hypothesize that some forum users would be more inclined to click on an invitation with an unfamiliar user name (“who is this?”) while other users might be more inclined to click on an invitation with a familiar name (“I remember you, you were interesting”).

We are not sure the degree to which this finding might generalize. It is possible that there are types of content or other domains for which familiarity algorithms are more useful. For example, in domains where the volume of traffic is much higher than ours, choosing familiar users could make the community feel more intimate, encouraging users to be social.

Research Question 3: Suggestions. Finding 5 and Finding 6 taken together show that in MovieLens, showing user names helped increase clicks, while showing movie titles did not. Why is this the case? One possible explanation is that showing the name of a poster emphasizes the social nature of the discussion forums. The perceived value of visiting the forums is enhanced by emphasizing features that are not available from the movie recommendation interface: the written opinions of other users.

CONCLUSION

We have investigated the usefulness of *personalized invitations*, an intervention designed to increase participation in an online discussion forum. We found that invitations had an immediate impact in MovieLens over the short-term, causing users to write and view more posts.

We experimented with several algorithms for choosing the content of invitations, with varying success. Our algorithm designs were influenced by theories from social psychology: uniqueness and familiarity. While uniqueness turned out to be a useful principle in our algorithm design, familiarity did not, perhaps confounded by users’ curiosity to learn about new things. It is possible that familiarity-based personalization algorithms would be more effective in other domains, especially larger communities where it is more difficult to locate content written by familiar users.

Invitations emphasizing the social nature of the discussion forum increased user activity, while invitations emphasizing other details of the forum were less successful. While showing the name of a movie did not lead to increased user interest, showing the name of a post author did. *Disagree*, the algorithm we consider the most successful in this work, emphasized the presence of social interaction and related the target user directly to a post author. The forums are a social space; users were more drawn to them when the social aspects were emphasized.

ACKNOWLEDGMENTS

Thanks to the MovieLens members who continue to visit our site and chat about movies. Thanks to members of GroupLens research for their feedback and support. This work is supported by a grant from the National Science Foundation (IIS-0324851).

REFERENCES

1. Bewley, W., Roberts, T., Schroit, D., and Verplank, W. Human factors testing in the design of Xerox's 8010 "Star" office workstation. *CHI 1983*.
2. Box, G. E. P., Hunter, W. G., and Hunter, J. S. *Statistics for Experimenters: An Introduction to Design, Data Analysis, and Model Building*. John Wiley & Sons, New York, 1978.
3. Bryk, A. S., and Raudenbush. *Hierarchical Linear Models: Applications and Data Analysis Methods*. Sage, CA, 1992.
4. Butler, B. When is a group not a group: An empirical examination of metaphors for online social structure. *Social and Decision Sciences*, 1999.
5. Chaiken, S., Liberman, A., and Eagly, A. H. Heuristic and systematic information processing within and beyond the persuasion context. In Uleman, J. S. and Bargh, J. A. (Eds.), *Unintended Thought*. Guilford Press, 1989.
6. Cosley, D., Frankowski, D., Terveen, L., and Riedl, J. Using intelligent task routing and contribution review to help communities build artifacts of lasting value. *CHI 2006*.
7. Drenner, S., Harper, M., Frankowski, D., Riedl, J., and Terveen, L. Insert Movie Reference Here: A System to Bridge Conversation and Item-Oriented Web Sites. CHI Notes, *CHI 2006*.
8. Festinger, L. A theory of social comparison processes. *Human Relations*, 7, 1954.
9. Greenspan, R. Surfers Prefer Personalization. <http://www.webcitation.org/5IsMR19Gt>. 2004.
10. Hill, W., Stead, L., Rosenstein, M., and Furnas, G. Recommending and evaluating choices in a virtual community of use. *CHI 1995*.
11. Lampe, C. and Johnston, E. Follow the (slash) dot: effects of feedback on new members in an online community. *Group 2005*.
12. Loewenstein, G. The psychology of curiosity: A review and reinterpretation. *Psychological Bulletin*, 116, 1994.
13. Ludford, P.J., Cosley, D., Frankowski, D., and Terveen, L. Think Different: Increasing Online Community Participation Using Uniqueness and Group Dissimilarity. *CHI 2004*.
14. McDonald, D. Internal and external evidence in the identification and semantic categorization of proper names. In *Corpus Processing For Lexical Acquisition*, 1996.
15. Nonnecke, B. and Preece, J. Lurker demographics: Counting the silent. *CHI 2000*.
16. Preece, J., Nonnecke, B., and Andrews, D. The top five reasons for lurking: improving community experiences for everyone. *Computers in Human Behavior*, 20, 2004.
17. Resnick, P., Iacovou, N., Suchak, M., Bergstrom, P., and Riedl, J. GroupLens: An open architecture for collaborative filtering of netnews. *CSCW 1994*.
18. Ridings, C. M., and Gefen, D. Virtual community attraction: Why people hang out online. *Journal of Computer Mediated Communication*, 10(1), 2004.
19. Schafer, J., Konstan, J., Riedl, J. E-Commerce Recommendation Applications. *Data Mining and Knowledge Discovery*, 2001.
20. Simon, H. Designing organizations for an information-rich world. In Greenberger, M. (Ed.) *Computers, communications and the public interest*. The Johns Hopkins Press, Baltimore, MD, 1971.
21. Webster A.S., Vassileva J. Visualizing Personal Relations in Online Communities. Workshop on Social Navigation and Community Based Adaptation Technologies, *Adaptive Hypermedia 2006*.
22. Zajonc, R. B. Attitudinal effect of mere exposure. *Journal of Personality and Social Psychology*, 9, Monograph Supplement, 1968.
23. Zipf, G.K. *Human behavior and the principle of least effort*. Addison-Wesley, Cambridge, MA, 1949.