

Think Different: Increasing Online Community Participation Using Uniqueness and Group Dissimilarity

Pamela J. Ludford, Dan Cosley, Dan Frankowski, Loren Terveen
University of Minnesota, Department of Computer Science and Engineering

4-192 EE/CS Building, 200 Union St. SE, Minneapolis, MN 55455, USA
[ludford|cosley|dfrankow|terveen]@cs.umn.edu

ABSTRACT

Online communities can help people form productive relationships. Unfortunately, this potential is not always fulfilled: many communities fail, and designers don't have a solid understanding of why. We know community activity begets activity. The trick, however, is to inspire participation in the first place. Social theories suggest methods to spark positive community participation. We carried out a field experiment that tested two such theories. We formed discussion communities around an existing movie recommendation web site, manipulating two factors: (1) *similarity*—we controlled how similar group members' movie ratings were; and (2) *uniqueness*—we told members how their movie ratings (with respect to a discussion topic) were unique within the group. Both factors positively influenced participation. The results offer a practical success story in applying social science theory to the design of online communities.

Categories & Subject Descriptors: H.5.3 [Information Interface and Presentation]: Group and Organization Interfaces-*Computer-supported cooperative work (CSCW)*.

Keywords: Online communities, recommender systems, social psychology, similarity, uniqueness.

INTRODUCTION

Robert Putnam has popularized the notion of social capital, the productive capacity that resides in social relations. Using data largely from the United States, Putnam has documented a significant decline in social capital over the past few decades [23]. We take this problem as a challenge, seeking ways to harness new technology like the Internet to create social capital. Online communities are not a simple substitute for bowling leagues or fraternal societies, but we believe they afford opportunities to develop alternative forms of productive social relations [26, 28].

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

CHI 2004, April 24–29, 2004, Vienna, Austria.

Copyright 2004 ACM 1-58113-702-8/04/0004...\$5.00.

A challenge to this vision is that many online communities fail to generate enough social capital even to sustain continued participation. For example, Butler [1] found 50% of social, hobby, and work mailing lists had no traffic over a 122 day period. Under-contribution is a problem even in communities that do survive: in a majority of active mailing lists, fewer than 50% of subscribers posted even a single message in a 4-month period [1]. Even in successful communities, questions can go unasked or unanswered.

A variety of communities therefore suffer from a deficit of visible content contribution. Although there are many ways to participate in online communities (including simply reading posts [22]), our study focuses specifically on overcoming problems caused by lack of visible content contribution. We do this by investigating new ways to spark generation of observable content.

We believe one reason online communities fail to elicit activity is social structures needed to sustain contribution have not been incorporated systematically into online community design and operation. One way to remedy this situation is to exploit insights from social theories that address why people contribute in face-to-face communities. We do that here, with the aim of developing innovative techniques and practical guidelines designers can use to increase positive contributions to online communities.

We report on the results of a field experiment using the MovieLens film recommender system [2] as a platform. The purpose of MovieLens is to take a set of movies a user had seen and rated, and return a new set of films he or she might enjoy based on the ratings provided. In the experiment, we expanded MovieLens' functionality, adding new online discussion groups with controlled design parameters. Specifically, we manipulated two factors that theories predict will influence participation in a group: *similarity* of group members and *uniqueness* of a member's qualification for a group task. We explain these concepts in detail in sections that follow.

In our experiment, similarity and uniqueness affected level of contribution. First, people in dissimilar groups contributed more. Second, people liked finding out how they were unique within a group, and providing them with such information increased their participation. We also

found the most active members of discussion groups were not those who had traditionally been most active in MovieLens: instead, the discussion community brought forth a new set of leaders. Finally, participants in the experiment contributed back to the larger community by rating significantly more movies over the course of the experiment than did a control group.

In the rest of the paper, we outline related work, describe the design and results of our experiment, and conclude by interpreting our results and identifying opportunities for future work.

RELATED WORK

Previous work has studied factors that lead individuals to participate in online communities. Hemetsberger and Pieters [9] studied open source software development, building on results from social psychological theory. They analyzed results from a survey on Slashdot.org that measured intrinsic and extrinsic reasons for participation. Key reasons people participated were to fulfill personal needs, to learn, and to advance the common good.

Several researchers have developed design principles for online communities that promote member participation [15, 16, 22]. Girgensohn and Lee [9] implemented a mix of the principles on two different web sites: CHIplace and Portkey. The sites offered features such as member profiles, individual activity level, and pages highlighting new members. Girgensohn and Lee found the design components were successful, but that a supportive infrastructure alone did not foster site usage. They identified new content and postings as necessary to promote ongoing user participation. Millen and Patterson investigated several potential influences on participation, including channeling community members to the same “place” to increase social density [18].

Nuances of facial expression, body language, and voice intonation are absent in online communication. Further, online communication may be anonymous, which affects the quality of interactions [8]. These properties may limit people’s ability to form substantive relationships online. Parks and Floyd [17] studied relationships formed in newsgroups and MUDs (Multi-User Dimensions). They found people do form substantive relationships online, more so in the richer, synchronous medium of a MUD than in asynchronous newsgroups.

Social psychologists have found people contribute less energy when they work with others than when they work alone [10]. A key cause for this decline is motivational: in a group, sometimes people think others will do the work. Karau and Williams’ collective effort model [14] tries to explain this “social loafing” and identify variables that reduce it. One such variable is whether individuals think their efforts affect the outcome of a task. One reason people contribute is they believe they bring unique skills or knowledge to the group. For example, if Susan is asked to

review ethnographic methods papers for a CHI conference, she is predicted to be more likely to agree (and to do a better job) if she is told that she has published more papers in this area than any other potential reviewers.

The influence of interpersonal similarity on relationship formation and group behavior has been much studied by social scientists. For example, studies have shown friends are likely to be quite similar to each other on a range of factors such as ethnicity, income, education level, religion, and profession [4, 6, 11]. Karau and Williams’ model also predicts that more similarity among members of a group will increase the individuals’ positive feelings toward other members and the group as a whole, and thus will increase their contributions. We apply this concept in our experiment, creating discussion groups of people who have either similar or dissimilar movie-related interests. Our purpose in doing so is to test the effect of similarity on participation and user satisfaction in online communities.

Like these research efforts, we seek to understand what makes people participate in online communities. However, our work is distinctive in several ways. First, we carried out a field experiment—we formed online groups in a controlled way that let us quantify the effect of various factors on participation. Second, we were directly guided by empirical results and theory—the factors of similarity and uniqueness we manipulated have firm grounding in the social psychological literature. Third, we conducted our research using a recommender system [24] as a platform—this let us implement the concepts of uniqueness and similarity quite naturally, yet we believe our implementations are fairly general.

EXPERIMENTAL PLATFORM

We performed our research in the context of the MovieLens film recommender. A recommender system collects a user’s opinions about items in a domain, e.g., their ratings of movies. It matches the user with others having similar taste, then recommends items these neighboring users rated highly, but that the target user has not rated.

Using the MovieLens recommender system for our experiment yielded several important benefits. First, MovieLens has over 70,000 registered users, over 2000 of who are active in a given month; thus we had a large potential user base for our experiment. Second, it let us draw on MovieLens data to calculate similarity and uniqueness. Similarity between users was computed by comparing their movie ratings. Uniqueness also was computed from ratings, and was relative to a particular topic and group of users. For example, if the topic were “What little-known film have you seen that you’d recommend to other people?” the system might find Frank was the only person in his group who had rated the movie “Imagine”, and furthermore, that very few members of the entire community had rated this movie. We relayed this information to Frank in case it helped him form a response

to the discussion topic. We further detail the uniqueness concept and computation process shortly.

EXPERIMENTAL DESIGN

Social science theories have identified many factors that influence individual contributions to a group, including: group size, group attractiveness, expectations of group performance, attractiveness of individual members of a group, the importance of one's contribution to group outcomes, explicit incentives, and the prospect of repeated interaction [4, 14, 17, 19, 25]. We considered investigating many of these factors before settling on similarity and uniqueness. Because these factors were well motivated theoretically, we could see how to implement them in the MovieLens context, and we also believed they could be applied to a broad range of online communities.

We studied these factors by creating a number of discussion forums in MovieLens. We implemented the forums as online message boards. The boards supported asynchronous conversation, which provided several benefits: users could participate at their convenience, long and thoughtful conversations were possible, and users didn't have to compete with each other to be "heard" as they might in synchronous chat.

Subjects

We recruited subjects by email, inviting MovieLens users who had rated at least 50 movies to participate in an experiment studying online communities. Of the nearly 8500 invitations sent, approximately 2800 bounced. 245 people volunteered to participate, resulting in a response rate of roughly 245/5700 or about 4%.

Subject ages ranged from 18 to 79 years, and the average age was 37. 27% were female, 73% male. Subjects were highly educated; 41% had earned a graduate degree, and 32% had a bachelor's degree. 91% of the subjects lived in the United States, 8% in Europe, and 1% in Asia. These demographics reflect those of MovieLens as a whole.

We asked subjects how often they participated in online discussion groups. 19% said they participated daily, 24% participated once a week, 28% participated between once a month and once every few months, and 29% had never participated in an online discussion forum prior to our experiment. In our analysis, we found prior forum experience had no effect on participation.

We provided a modest incentive to encourage MovieLens users to participate in the study. Subjects who participated at a minimum level (posting at least one message during four of the five weeks of the study) were eligible for a drawing in which five Amazon.com gift certificates were given away. Our analysis discounted the effect of the incentive by analyzing participation beyond the requested minimum where appropriate.

Controlling the Experimental Factors

The experiment used a 2 x 2 design; we formed 8 groups (2 in each experimental condition) consisting of 28 or 29 people¹. All participants were expressly recruited for the experiment via the email process described earlier. Four of the groups consisted of members with similar ratings of movies; four consisted of members with less similar ratings. Four of the groups received weekly email messages advising them of a unique perspective they could bring to the current discussion topic. The other four groups acted as control groups for the uniqueness condition. They received weekly email messages; these messages, however, simply announced the new discussion topic.

The Experimental Task

The experimental task was simple: subjects discussed movies using the online forum. The experiment ran for five weeks, from the end of July to the start of September 2003.

Subjects accessed their forum from the MovieLens home page. A link to the forum was visible only to experimental subjects. Subjects could only view and participate in the forum they were assigned to.

Subjects could talk about whatever they wanted in the forums. In addition, the research team posted one discussion topic each week. Our posts served multiple purposes. They kept the forums from being empty when the experiment began and assured new content appeared regularly. For subjects in the uniqueness condition, we computed an exclusive movie rating history with respect to these topics. The five weekly topics were:

1. What is a little-known film you have seen that you'd recommend to other people?
2. What acting performance was worthy of an Oscar, but did not win?
3. Was acting better in the 1950s (and earlier) than it is today?
4. Take a look at the films MovieLens recommends for you. There's probably at least one on the list you're unfamiliar with or are not sure you want to see. Discuss a recommendation you're unsure about with the group.
5. In your opinion, what makes a love story 'click' in a movie?

The Uniqueness Condition

Subjects in the 'uniqueness' condition received a weekly message telling them how their MovieLens ratings differed from others in their group relative to a discussion topic. For example, consider the first topic we posted: "What is a little-known film you have seen that you'd recommend to other people?" For this topic, we found movies for each subject that:

¹ We also assigned 15 people to a beta group who went through the process one week ahead of the 8 main groups.

- had been rated by fewer than 1000 MovieLens users (thus, they were little-known),
- the subject had rated favorably (thus, the subject was likely to recommend it), and
- no one else (or few others) in the subject’s group had rated (thus, the subject had a unique perspective).

We did not explicitly tell subjects to mention their uniqueness information when they posted; instead, we simply explained they might find the information relevant to the discussion topic.

Calculating Uniqueness

The movies we considered “unique” for a given subject were a function of the discussion topic, the subject’s movie ratings, and the discussion group’s movie ratings. To compute the set of unique movies, we first generated a list of candidates relevant to the discussion topic. For example, for the topic, “What makes a love story ‘click’ in a movie?” candidate movies were romance films the subject had rated.

We then ordered the set of candidate movies by distinctiveness. We first sorted movies by how often they had been rated by others in the group: movies rated by fewer group members were deemed more unusual and thus were preferable. For movies that tied according to this criterion, we used the subject’s movie rating as a tiebreaker. We did additional filtering and sorting each week to assure the information related to the discussion topic as closely as possible. Sometimes the most natural definition of “relevant” included too few movies, making it impossible to calculate uniqueness information for some subjects. In those cases we had to relax our criteria.

The Similarity Condition

While the uniqueness condition was implemented via weekly email, similarity was fixed at group creation time. “Similarity” for two users means they tend to see the same movies and agree on their evaluations of the movies. “Dissimilar” users either see different movies or disagree on whether they like movies they have both seen. We formed four groups of similar users and four of dissimilar users. Forming groups required two things: choosing a similarity metric and clustering the users.

Choosing a metric. Collaborative filtering systems consider two users similar when their ratings tend to agree. Pearson correlation and cosine similarity are two popular methods for computing agreement between sets of ratings [27]. However, Pearson correlation considers only items both users have rated. This can lead to spurious correlations between users who have rated only a few common items. Cosine similarity uses all of the users’ ratings, so is less likely to report spurious correlations. Both metrics express similarity as a number between -1 and 1, which we found hard to interpret: are two users with a similarity of 0.29 much more similar than a pair with a similarity of 0.15?

We therefore chose *co-agreement* as our similarity metric. We used 3 stars (on a 0.5 to 5 star scale) as a threshold, and said that two users agreed on a movie when both rated the

movie below the threshold or both rated it at or above the threshold—intuitively, both users either disliked or liked the movie. The co-agreement similarity score between two users is the number of movies they agreed upon.

This metric has several nice properties. It is fast to compute. It correlates well with cosine similarity ($r^2 = 0.97$ on our data). Most important, we think it expresses the difference between similar and dissimilar groups in a way that is easy to understand. Table 1 summarizes the average pair-wise similarity between members of similar groups and between members of dissimilar groups.

	Co-agreement	Cosine	Pearson
Similar	102 movies	0.28	0.16
Dissimilar	57 movies	0.16	0.11

Table 1: Average pair-wise similarity between members of similar and dissimilar groups.

Clustering the users. A standard clustering algorithm was not appropriate for us since we wanted to create both clusters and “anti-clusters”. Our algorithm builds similar (dissimilar) groups by starting with the most (least) similar pair of users, then adding the user that resulted in the highest (lowest) average pair-wise similarity among group members. The algorithm adds users to a group until it has the desired number of members. After enough groups are formed, it improves the results by swapping users between groups as long as the total difference between similar and dissimilar groups increases. We imposed one additional constraint, forcing the algorithm to assign subjects to groups so that the distribution of number of ratings among members in each group was roughly the same.

Hypotheses

We are now in a position to state the hypotheses we studied in the experiment. The first three concern **uniqueness**:

- H1:** *People contribute more to online communities when shown personalized uniqueness information.*
- H2:** *People given uniqueness information exploit this information when participating in the community.*
- H3:** *People like being shown uniqueness information.*

We also stated two hypotheses about **similarity**:

- H4:** *People contribute more to online communities when they are in a discussion group with others similar to themselves.*
- H5:** *People can identify similar people through online community interaction.*

Finally, we stated one additional hypothesis:

- H6:** *Users who are active MovieLens raters will be most active in the discussion groups.*

Measures

The 230 subjects in 8 forums posted a total of 1473 messages over the course of the study. 163 subjects posted

at least one message. Posting followed an inverse power law: 9% of the subjects accounted for 50% of all posts.

We tested our hypotheses concerning participation by analyzing the quantity of posts in the various conditions and the results of a post-experiment survey.

RESULTS

The key finding was that dissimilar groups who were given uniqueness information were the most active communities throughout the experiment.




Figure 1: Number of posts by condition and week.

Unique groups posted more messages than no-unique groups. In total, unique groups averaged 49 posts per week, while no-unique groups averaged 21 posts. The difference was significant using a t-test ($p < .01$, $t = 3.24$, $df = 38$). There was a spike of activity in week 1 (see Figure 1), so we ran a second test to determine if differences in posting still were significant once behavior stabilized in weeks 2-5. We also did a third test to discount the effect of the participation incentive, this time counting only the number of posts in excess of the minimum requirement during weeks 2-5. In both cases, the differences remained significant (see Table 2 for details). These results support hypothesis **H1**: ‘People contribute more to online communities when shown personalized uniqueness information’.

	Unique	No-Unique	T-test Detail
All-Inclusive	49	21	$p < .01$, $t = 3.24$, $df = 38$
Weeks 2-5	44	18	$p < .01$, $t = 2.77$, $df = 30$
Above Incentive	16	8	$p < .01$, $t = 2.07$, $df = 106$
	Similar	Dissimilar	T-test Detail
All-Inclusive	20	51	$p < .01$, $t = 3.67$, $df = 38$
Weeks 2-5	17	45	$p < .01$, $t = 3.13$, $df = 30$
Above Incentive	7	17	$p < .02$, $t = 2.54$, $df = 106$

Table 2: Posting behavior in unique vs. no-unique and similar vs. dissimilar groups.

We conducted the same 3 tests to measure whether similar groups posted more than dissimilar groups did. The results were the opposite of what we expected: in each test, dissimilar groups posted more than similar groups, and the differences were statistically significant. Table 2 presents the details. These results were contrary to our hypothesis

H4: *people contribute more to online communities when they are in a discussion group with others similar to themselves.* We will discuss possible interpretations and implications of this unexpected finding shortly.

In addition to testing effects of uniqueness and similarity on community participation independently, we checked whether the two factors interacted to boost participation. We found they did. We did a three-way ANOVA test comparing the number of posts made in each group during each week of the experiment. The ANOVA factors were similarity condition, uniqueness condition, and week. We contrasted all four types of groups (similar/no-unique, similar/unique, dissimilar/no-unique, dissimilar/ unique) and found an interaction between the uniqueness and similarity condition ($p < .05$, $df = 32$). However, we note it was the combination of uniqueness and *dissimilarity* that caused participation to increase. We can only speculate why these factors interact, and believe further study is needed to understand the results.

Hypothesis **H2** postulated *people given uniqueness information would exploit this information in contributing to the community.* Specifically, we expected subjects to use the uniqueness information we sent them in their posts, e.g., we expected them to discuss the movies we mentioned. We examined responses to post survey questions to test the hypothesis and found mixed results. 44% percent of subjects said they both responded to the related topic and used the uniqueness information in their post, while the remaining 55% said they responded to the question but did not use the uniqueness information. To follow up, we asked subjects whether they found the uniqueness information relevant to the discussion topics. A linear regression test showed a strong correlation between whether subjects found the uniqueness information relevant and whether they used it in the discussion forum ($r^2 = .95$, $p < .05$). Since subjects did not unequivocally use uniqueness information in their posts, the results do not confirm the hypothesis. Nonetheless, the results are interesting: they imply that improving the uniqueness algorithm to increase relevance would make its suggestions more useful.

H3 stated *people like being shown uniqueness information,* and data from the post-survey supported the hypothesis: 82% of the subjects who received uniqueness information said they benefited by receiving it. In a general comments section, a handful of subjects said they found the uniqueness information thought provoking, and mentioned they liked learning how their movie rating history differed from others in the discussion group. For example, (referring to the uniqueness email) one subject said, “...it was a neat reminder of how I (had) rated some movies that I hadn't thought about in a while. It was also nice to see what movies I rated that very few others had...”

H5 hypothesized *people can identify similar people through online community interaction.* Again, answers to the post-survey supported this hypothesis. We gave subjects a list of

group members who had posted at least once and asked how similar their views were to individuals on the list. People in similar groups rated their views similar to other members more often than people in dissimilar groups did. The responses were measured on a 4-point likert scale (from ‘Very Similar’ (1) to ‘Very Dissimilar’ (4)), and were significantly different using a t-test ($p < .01$, $t = 4.24$, $df = 860$); Table 3 presents details. Second, we asked subjects whether other members of their group generally disagreed with their views. Members of dissimilar groups agreed with this statement more often than members of similar groups did. The difference was significant using a t-test ($p < .01$, $t = 3.34$, $df = 105$).

To illustrate the different interaction styles in the ‘similar’ and ‘dissimilar’ conditions, we provide excerpts from forums below. The discussion topic was: “Was acting better in the 1950s (and earlier) than it is today?”

Similar group:

Hugo: “Similar to several of you, I don't think acting is better or worse for that matter. The points about technology, special effects, and writing are quite valid...”

Ranger: “I have to agree with the general consensus that today's acting is no worse than that in the 50s...”

Dissimilar group:

Wibby: “(Acting in the 50s) was different. And this is partly due to the changes in the way directors make movies...”

Jake: “I'm not sure what you are thinking of here because I can hardly agree with you. I would take the exact opposite view on most of your points and will explain why...”

	Similar Condition	Dissimilar Condition
Perceived similarity with individuals in group	2.29	2.54
Perceived disagreement level with group	2.12	2.71

Table 3. Subject’s assessment of individual group member similarity and group disagreement level.

Finally, **H6** predicted *people who were active in the MovieLens community would be the most active participants in the discussion groups*. We compared the number of movies subjects had rated with the number of posts they made. There was no correlation, so this hypothesis is not confirmed. The result is nonetheless interesting: it suggests discussion groups attract and motivate a different type of user than recommender systems in general.

We found a noteworthy converse relationship: people who posted to the experimental discussion groups rated an average of 65 movies during the experiment. A control group of MovieLens users who had been invited to participate but were not in the experiment rated an average of 30 movies during the same period. The difference in number of ratings was significant using a t-test ($p < .01$,

$t = 4.41$, $df = 861$). So, being active in the discussion forum also led people to contribute more to the community in other ways: more ratings means better recommendations for all members of the MovieLens community.

Table 4 summarizes our results. People contributed more to a discussion when they were told something unique they could bring to bear and when they were in groups with others with dissimilar views. Also, number of posts did not correlate with number of movie ratings on the recommender side of MovieLens. Instead, a new group of users were most active in the movie discussion forums.

Hypothesis	Supported
H1: <i>People contribute more to online communities when shown personalized uniqueness information.</i>	Yes
H2: <i>People given uniqueness information exploit this information when participating in the community.</i>	Mixed Results
H3: <i>People like being shown uniqueness information.</i>	Yes
H4: <i>People contribute more to online communities when they are in a discussion group with others similar to themselves.</i>	No
H5: <i>People can identify similar people through online community interaction.</i>	Yes
H6: <i>Users who were active MovieLens raters were most active in the discussion groups.</i>	No

Table 4. Summary results for our hypotheses.

DISCUSSION

We now turn to how designers of and researchers in online communities might benefit from our results.

Uniqueness

Community members liked receiving information about the unique perspective they brought to the group (**H3**) and participated more because of it (**H1**). This suggests a design guideline: “To increase members’ contributions to and satisfaction with a community, tell them how they are special with respect to the group and its purpose.”

Implementing this guideline requires developers have data about group members and can easily process it to discover how individuals differ from each other. In our study, we mined MovieLens data to find how subjects’ movie ratings differed from others in the group. However, applying the uniqueness concept does not require a recommender system, or even an existing database of user information.

For example, the GardenWeb site currently has a forum called ‘Name that Plant’. Gardeners submit pictures of plants and hope other community members will help identify them. To implement the uniqueness concept, the site could track the discussions and posts gardeners view, thus inferring information about their locale and specialty gardening knowledge. As new ‘Name that Plant’ requests appear, GardenWeb could use information retrieval

techniques to match the gardener's location and/or knowledge with requests as relevant. When users log in, GardenWeb could tell them if they are in a unique position to identify a plant and point them to the appropriate post.

Making uniqueness easy to compute is crucial, as many communities lack resources to maintain infrastructure [8]. If it's too much work to produce uniqueness information, community maintainers may not bother. We consider two approaches to minimize the workload.

One approach is to use automation, e.g., GardenWeb's proposed use of information retrieval techniques. The uniqueness calculations were largely automatic in our study, too. However, we generated the list of relevant movies for each topic by manually crafting database queries, which was too much effort. What if we could have used our members to help the system choose relevant movies? In the future, MovieLens could ask topic creators to provide a few movie titles relevant to the topic and use the recommender systems' knowledge of similarity between movies to choose a plausible set of relevant movies. Users who have rated the relevant movies could be advised of the post and their unique perspective for contribution.

A second approach to generating uniqueness information is to distribute the responsibility to community members. GardenWeb could request users create profiles, directly providing information about their locale and specialty gardening knowledge. In addition, people submitting 'Name that Plant' requests could supply keywords potentially relevant to the unknown plant. The system could use the keywords to find members who might be in a unique position to identify the plant. Whatever the approach, developing structures for self-maintaining communities is a research area with the potential for major real-world impact, and one deserving additional study.

Finally, the success of uniqueness information in increasing contributions and users' satisfaction suggests a number of related questions. For instance:

- How unique does information have to be?
- How do people respond to information showing what they have *in common* with a group?
- Does the effect of receiving uniqueness information deteriorate over time?

Similarity

Contrary to our expectations, those in similar groups contributed *less* than those in dissimilar groups (**H4**). We consider why the results countered our prediction, suggest how the outcome could be applied in practice, and outline ideas for additional research using similarity.

To review, higher participation in dissimilar groups surprised us because social theory suggests people are attracted to others similar to themselves. We therefore anticipated more participation in similar groups. Informal content analysis suggests a different picture: people with dissimilar views had longer exchanges, appearing to banter

back and forth defending their positions. Preece has found members of certain communities seek support and reassurance rather than friction [21], so why then did subjects in our experiment participate despite group discord? Recent work by Guerin suggests people are more drawn to controversy than cooperation in conversation [10]. In communication, participants seek to hold the attention of the other conversants, and Guerin contends scandalous or controversial topics are effective tools for doing so.

In our study, all MovieLens users share a passion for movies. Thus, our similar groups may have lacked enough difference of opinion to spark discussion. Further, we note subjects in all conditions overwhelmingly reported other group members respected their opinions. When disagreement occurred, it was polite disagreement, not rancorous discord. In fact, one individual in a dissimilar group commented to another, "We may disagree, but that just makes for good discussion." We recognize some online groups are dominated by flame wars between people who disagree with each other, and such exchanges shed more heat than light. When disagreement is permissible, however, it promises to encourage discussion. We propose the following guideline: "In order to encourage participation, favor creating dissimilar groups in situations where disagreement can be tolerated."

To use either similarity or dissimilarity to create groups, web site developers might choose to create multiple parallel communities as we did. When users first enter a web site, they might not see any community features. Once the system learns enough about users, it can place them in communities they are most likely to enjoy and contribute to. As a bonus, placing people in smaller communities makes it easier to determine their unique contributions. It also would be possible to make user profiles visible, leaving it up to individuals to decide whether to affiliate with others having congruent or differing views.

Opportunities await researchers who can identify and exploit other kinds of similarity besides similarity of taste. Cosley, Ludford, and Terveen [3] showed demographic similarities affect how well strangers get along online. Further, studies in the sociology of friendship [13] have shown that the context in which a friendship is formed influences the kinds of similarity that matter. Further research is needed to determine which types of similarity matter most in online communities.

Other factors

Uniqueness and similarity are two of many factors social theories suggest will affect user contributions to a community. For example, the size of an online community clearly is an important factor. Our post survey asked whether there were enough people in the discussion group; the majority of subjects said the groups should have had more participants. Finding creative ways to use group size and other factors to increase the success of online communities is a rich area for future research.

CONCLUSION

This work illustrates a promising approach to the design of online communities. We identified factors that social psychology research says should affect participation and found ways to implement them in a field experiment. Our specific results are strong—forming groups with diverse perspectives and showing people their unique qualities relative to a topic increase contributions. More generally, this study illustrates a research program with the potential to strengthen scientific foundations and increase the practical success of online communities.

ACKNOWLEDGEMENTS

We gratefully acknowledge Greg Grossmeier and Mike Cassano for their efforts in helping choose and implement the community software. This research was supported by NSF Grant IIS-0324851 and a University of Minnesota Grant-In-Aid.

REFERENCES

1. Butler, B. Membership Size, Communication Activity, and Sustainability: A Resource-Based Model of Online Social Structures. *Information Systems Research* 12, 4 (2001), 346-362.
2. Cosley, D., Lam, S.K., Albert, I., Konstan, J., & Riedl, J. Is Seeing Believing? How Recommender Systems Influence Users' Opinions. *Proc. CHI 2003*, 585-592.
3. Cosley, D., Ludford, P., & Terveen, L. Studying the Effect of Similarity in Online Task-Focused Interactions. *Proc. GROUP 2003*, 321-329.
4. Fehr, E. & Schmidt, K. A Theory of Fairness, Competition and Cooperation. *Quarterly Journal of Economics* 114 (1999), 817-868.
5. Fischer, C. *To Dwell Among Friends: Personal Networks in Town and City*. University of Chicago Press, 1982.
6. Fischer, C., et al. *Networks and Places: Social Relations in the Urban Setting*. New York: Free Press, 1977.
7. Fishkin, J.S. *The Voice of the People*. Yale University Press, 1997.
8. Friedman, E. & Resnick, P. The Social Cost of Cheap Pseudonyms. *Journal of Economics and Management Strategy* 10, 2 (2001), 173-199.
9. Girgensohn, A., & Lee, A. Making Web Sites Be Places for Social Interaction. *Proc. CSCW 2002*, 136-145.
10. Guerin, B. Language use as social strategy: A review and an analytic framework for the social sciences. *Review of General Psychology* 7, 3 (2003), 251-298.
11. Hemetsberger, A., & Pieters, R. Fostering cooperation on the Internet: social exchange processes in innovative virtual consumer communities. *Proc. Association for Consumer Research* (2001).
12. Ingham, A.G., Levinger, G., Graves, J., & Peckham, V. The Ringelmann effect: Studies of group size and group performance. *Journal of Experimental Social Psychology* 10, 4 (1974), 371-384.
13. Jackson, R.M. Social Structure and Process in Friendship Choice. In *Fischer et al.: Networks and Places*. New York: Free Press, 1977, 59-78.
14. Karau, S.J., & Williams, K.D. Social loafing: A meta-analytic review and theoretical integration. *Journal of Personality & Social Psychology* 65, 4 (1993), 681-706.
15. Kim, A.J. *Community Building on the Web*. Peachpit Press, Berkeley, CA, 2000.
16. Kollock, P. Design Principles for Online Communities. *Proc. Harvard Conf. on the Internet and Society* (1996).
17. Kraut, R.E. Applying Social Psychological Theory to the Problems of Group Work. In J. Carroll (ed): *HCI Models, Theories, and Frameworks*. 2003, 325-356.
18. Millen, D.R., & Patterson, J.F. Stimulating Social Engagement in a Community Network. *Proc CSCW 2002*, 306-313.
19. Oliver, P.E., & Marwell, G. The Paradox of Group Size in Collective Action: A Theory of the Critical Mass. II. *American Sociological Review* 53, 1 (1998), 108.
20. Parks, M., & Floyd, K. Making Friends in Cyberspace. *Journal of Communication* 46, 1 (1996), 80-97.
21. Preece, J. Empathic communities: reaching out across the Web. *Interactions* 5, 2 (1998), 32-43.
22. Preece, J. *Online Communities: Designing Usability, Supporting Sociability*. John Wiley and Sons, Ltd., England, 2000.
23. Putnam, R. *Bowling Alone*. Simon & Schuster, 2000.
24. Resnick, P., & Varian, H.R., guest editors, *CACM, Special issue on Recommender Systems* 40, 3 (1997).
25. Rabin, M. Incorporating Fairness into Game Theory and Economics. *American Economic Review* 83 (1993), 1281-1302.
26. Resnick, P. Beyond Bowling Together: SocioTechnical Capital. In J. Carroll (ed): *HCI in the New Millenium*. Addison-Wesley, 2002, 247-272.
27. Sarwar, B.M., Karypis, G., Konstan, J.A., & Riedl, J. Item-based collaborative filtering recommendation algorithms. *Proc. WWW 2001*, 285-295.
28. Sproull, L., & Kiesler, S. Transforming Public Volunteer Work. *Proc. Intl. Conf. on the Economic and Social Implications of Information Technology* (2003).
29. Van Alstyne, M. & Brynjolfsson, E. Electronic Communities: Global Village or Cyberbalkans? *Proc. 17th International Conf. on Information Systems* (1996).
30. Verbrugge, L.M. The Structure of Adult Friendship Choices, *Social Forces* 56, 2 (1977).