User Motivations and Incentive Structures in an Online Recommender System

F. Maxwell Harper[1], Joseph A. Konstan[1], Xin Li[2], Yan Chen[2]
University of Minnesota, USA[1], University of Michigan, USA[2]
{harper, konstan}@cs.umn.edu, {xinl, yanchen}@umich.edu

Abstract. Ratings-based recommender systems are one type of online community that relies on user contributions. We present an overview of the implicit incentive structures that motivate rating behavior in one such system, MovieLens. We conducted a survey of MovieLens users to determine their motivations, and formalized these findings in a parameterized economic model. We found that users are motivated to rate by different factors. Some rate to improve their recommendations, others rate because it is fun. We are currently investigating the effects of introducing explicit incentives in the form of personalized messages that compare users with one another.

Introduction

Ratings-based recommender systems depend on user contributions. Many of these systems avoid domain-specific rules and algorithms by using ratings data alone to generate personalized recommendations. If users do not contribute ratings to the community, especially for new and rarely-rated items, these systems lose their ability to produce recommendations–their main purpose for existence.

In the absence of an explicit incentive structure, what motivates users to rate? We have been conducting research to understand user motivations to contribute ratings in MovieLens, an online movie recommender system. In this paper, we present some of our findings relevant to the study of incentive structures in online communities. In the first section, we give an overview of the design of MovieLens, along with some discussion of how this design affects participation. We then present results from a survey of user motivations and the resulting parameterized economic model that we built to formalize our understanding. In the second section, we present our ongoing research into the effects of displaying messages that compare users to one another.

Modeling User Motivations in MovieLens

MovieLens does not contain explicit incentives to motivate contributions. Yet, many users continue to rate movies well after it is personally beneficial to do so from the perspective of the recommendation algorithm. What are the properties of MovieLens that induce this behavior? In this section, we describe the MovieLens system in terms of its purpose and user population, and we present the results of building an economic model of user rating behavior.

MovieLens Incentives

MovieLens provides interfaces for searching for movies and for rating movies on a 1-5 star scale. Any movie that a user has not rated will have a predicted rating associated with it (see [Sarwar 2001] for an overview of the recommender algorithm). Figure 1 shows a typical MovieLens screen displaying several movies along with their predicted ratings. The MovieLens database contains about 8,800 movies, 98,000 registered users, and 12.1 million ratings as of September, 2005.

While MovieLens does not currently use explicit incentives to reward ratings, there are implicit structures. New users are given incentives to build an initial profile. When users first sign up, they are given lists of movies to rate (see [Rashid 2002] for details), and told that they will begin to receive personalized recommendations once they have reached 15 ratings. We effectively withhold the reward of movie recommendations until the user has told us enough about his or her movie preferences. About 35% of the people who sign up for MovieLens do not complete the initial 15 ratings, which may indicate that the cost is too high relative to the perceived benefit for these users.

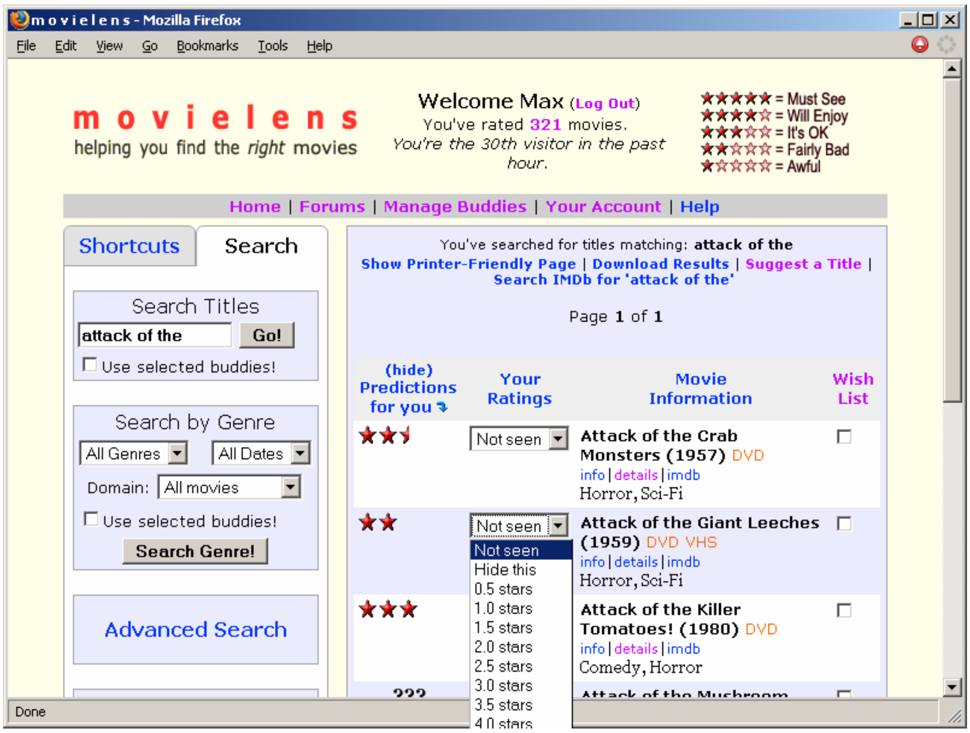

Figure 1: The MovieLens Interface

Users are also told that rating more movies will help them receive more accurate predictions. In this way, MovieLens continues to entice users with the promise of better movie recommendations in exchange for more ratings. However, many MovieLens users who have effectively maximized their recommendation quality continue to rate movies. Either they do not perceive the decaying advantage of providing more ratings, or there are other factors which motivate them to rate.

User Survey on Motivations

To better understand why users behave the way that they do in MovieLens, we conducted a survey of 358 users in June, 2004[1]. We solicited users by posting a banner message on the MovieLens home page. Only users with 30 or more ratings and 3 or more logins to the system were invited to participate. Some relevant results from the survey are summarized in the following list:

- **Reasons to use MovieLens**. Users were asked to rank their top-three reasons for using MovieLens. Not surprisingly, viewing movie recommendations was the most popular response, chosen as a top-three reason by 90% of the respondents. However, the second most popular reason to use MovieLens was to rate movies, chosen as a top-three reason

[1] See http://www-users.cs.umn.edu/~harper/group2005.html for screenshots of the survey and full results.

by 70% of the respondents, indicating that for at least some users, rating is not a means to an end, but an end of its own.

- **Reasons to Rate Movies**. Users were asked to rank their top-three reasons for rating movies. 87% of the respondents answered that improving their recommendations was a top-three reason to rate. More interestingly, the second and third most popular reasons to rate were to keep a list of movies (chosen by 54%), and because rating is fun (chosen by 48%). This provides further evidence that the fun of rating is one explanation as to why so many users continue to rate movies long after it has stopped benefiting their recommendation quality. In general, few users claimed to rate movies to exert influence on others. This is perhaps due to the lack of visible social queues in the recommendation interface.
- **Changing Experiences Over Time**. Users were asked to compare their experiences as a new user to their current experiences. 73% of the respondents agreed that their movie recommendations had become more valuable since they were a new user, and 86% agreed that rating movies had improved their recommendations as a new user. However, only 40% of the respondents agreed that rating movies continues to help improve their recommendations.

The survey results confirmed our beliefs that users rate to get better movie recommendations, and that the effect of getting better recommendations diminishes over time. However, the results reveal that the implicit fun of rating movies and list-keeping is a powerful motivator for users to contribute to the system. As a result, we believe that emphasizing the fun aspects of rating or building interfaces that encourage list-keeping would increase the volume of contributions. While building interfaces that emphasize users' influence on others may have some effect, it is likely that we would first have to emphasize in the MovieLens interface the social nature of our recommendation algorithm.

An Economic Model of User Rating Behavior

To formalize our findings about our users, we built a parameterized economic model of user rating behavior in MovieLens [Harper 2005]. The model was based on both survey data and behavioral data. The independent variables in the model included indicators designed to estimate a user's marginal costs and benefits (monetized!), how unique that user's taste in movies is, and how much fun that user has using the system. There was much trial and error in building the model; after all was said and done, we were able to account for 34% of the variability of user rating behavior, a solid result in the world of economic modeling of human behavior.

We can use this model to start thinking about personalized interfaces designed to increase user contributions. Now that we can efficiently fit users into this model, we can choose to emphasize different elements of the system to them. Users who most directly benefit from prediction quality can be given updated information on the quality of their recent predictions and the estimated increase in quality from the next quantum of ratings. Users who are more interested in the fun of rating itself can receive different cues and prompts.

Comparing Users With One Another

Subsequently, we've used our economic model to investigate questions about the effectiveness of different types of messages in soliciting contributions. Specifically, we are interested in understanding the effect of sending messages to users comparing their performance to that of other users. When users are told that they are above or below average, how do they respond? Will users who are told they are below average pick up the pace, or leave in frustration? Will users who are told they are above average be more willing to perform actions beneficial for the community?

Theories from the social sciences provide some predictions concerning these questions (for example, Inequality Aversion theory [Fehr 1999]). At a very high level, these theories assert that people tend to conform to the most prevalent behavior in their community and that some people are willing to sacrifice some amount of personal gain to increase equality across the community. Controlling for factors such as an individual's tendency towards altruism, we would expect to observe below-average individuals to be more likely to perform personally beneficial (i.e. selfish) actions. In MovieLens, we might expect below-average users to respond to a comparison by rating popular movies, ratings that provide little value to the community, but that may help the user receive more accurate predictions. Conversely, we would expect above-average users to be more likely to perform actions beneficial for the community, such as rating rarely-seen movies or maintaining the MovieLens database.

To test these theories, we conducted an online, controlled experiment in the summer of 2005. We randomly divided 398 email-recruited users into three experimental groups. Subjects in the first group received a control email newsletter containing some information about their MovieLens profile. Subjects in the second group received an email newsletter comparing the number of movies they had rated with the average number of ratings for similar users. Subjects in the third group received an email newsletter comparing their *net benefit* in MovieLens with that of similar users.[2] All of the newsletters contained the same five links to MovieLens; the non-control groups were told how their actions would affect their standing compared with other users. We measured how often the different links were clicked, as well as what actions users actually took

[2] See http://www-users.cs.umn.edu/~harper/group2005.html for sample email newsletters.

in MovieLens. We conducted pre- and post-experiment surveys to assess user motivations, calculate net benefit scores, and assess the degree of altruism and other personal characteristics.

We are currently in the process of analyzing the results of this experiment and will present our findings at the workshop. We expect to find that messages explicitly comparing users will provide a strong stimulus for motivating contributions. Understanding the effects of messages about users' relative standing will allow designers to better understand when and how prominently to display messages comparing users with one another in order to increase contributions.

Conclusions

We found that MovieLens users contributed ratings not just to improve their recommendation quality, but also because the rating activity is fun. This result points to the fact that fun is an implicit incentive that may lead users to increased levels of contributions. Ideally, we would recognize the personal motivations for each user and personalize the interface accordingly. Our economic model was an early step towards this goal. Our recent work has explored the effects of more explicit incentives: messages comparing users with one another. We will report the results of this study at the workshop.

Acknowledgments

This work is supported by a grant from the National Science Foundation (IIS-0324851).

References

Fehr, E. and Schmidt, K. (2000). A theory of fairness, competition, and cooperation. *The Quarterly Journal of Economics*, 114(3), 817-868.

Harper, F., Li, X., Chen, Y., Konstan, J. (2005). An Economic Model of User Rating in an Online Recommender System. In *Proceedings of The 10th International Conference on User Modeling*, Edinburgh, UK.

Rashid, A., Albert, I., Cosley, D., Lam, S., McNee, S., Konstan, J., & Riedl, J. (2002). Getting to Know You: Learning New User Preferences in Recommender Systems. In *Proceedings of the 2002 International Conference on Intelligent User Interfaces*, San Francisco, CA, pp. 127-134.

Sarwar, B., Karypis, G., Konstan, J., and Riedl, J. (2001). Item-based collaborative filtering recommendation algorithms. In *Proceedings of the 10th International World Wide Web Conference*, Hong Kong.