Cognitive Heuristics and Collective Opinions in Peer Recommendation

KRISTINA LERMAN, University of Southern California
TAD HOGG, Institute for Molecular Manufacturing

1. INTRODUCTION

The world presents far more information than people have the capacity to examine. As a result, humans have evolved to use cognitive heuristics to decide quickly what information to pay attention to. For example, people pay more attention to items near the top of a list than those below [Payne 1951]. A consequence of this cognitive heuristic, called *position bias*, is the strong effect the presentation order — item ranking — has on individual choices. It affects which items in a list of search results users click on [Buscher et al. 2009], and the answer they select in response to multiple choice questions [Blunch 1984]. Another common cognitive heuristic is *social influence*: people pay attention to the choices of others. Social influence affects most daily decisions, such as what to buy and who to vote for. Studies showed that social influence biases individual judgements [Salganik et al. 2006; Muchnik et al. 2013], creating an "irrational herding" effect that can obscure the underlying quality of choices.

Cognitive heuristics also play an important role online, where rapid proliferation of user-generated content makes it difficult to identify high-quality items. Since people often do not have time or energy to evaluate all available choices, they may rely on the opinions of others. Crowdsourcing, peer recommendation, and markets are some of the mechanisms for aggregating individual decisions into a collective opinion, which can help individuals identify high-quality items. The choices content providers make about how and what information to display to users has a profound effect on collective behavior. For instance, the choice of how to rank items shown to users can significantly affect outcomes of peer recommendation, since people will pay more attention to items near the top of the list. Similarly, displaying a social signal, which shows people how many others have liked an item, will also affect how much attention it receives. These effects combine with item quality to determine collective outcomes of peer recommendation.

We disentangle some of these effects through controlled web-based experiments using Amazon Mechanical Turk. Our studies quantify how position bias and social influence bias affect individual choices and collective outcomes of peer recommendation, and clarify results of earlier studies. For example, we demonstrate that large variation in popularity can arise in the absence of social influence simply due the attention items receive due to their position [Lerman and Hogg 2014]. In addition, we measure how social influence signals affect popularity after controlling for item quality and position. We find that social influence causes people to rely on others, rather than personal judgement, to determine whether the item is interesting. Although this contributes to the "irrational herding effect" [Salganik et al. 2006; Muchnik et al. 2013], social influence has a benefit that was not appreciated previously: by reducing the effort required to evaluate items, it increases the efficiency of peer recommendation.

We also demonstrate that we can leverage people's innate cognitive biases to more effectively aggregate collective opinions in peer recommendation. We show that simply by changing item ranking, we can distribute collective attention more homogeneously over higher quality items, reducing the unpredictability and variance of peer recommendation outcomes.

2. EXPERIMENTAL DESIGN AND RESULTS

Using Amazon Mechanical Turk (Mturk) we experimentally evaluated the impact of item ranking policies and social signals on peer recommendation. Experiments provide a controlled setting to isolate the effects of list position and social influence from the items' quality. In our experiments, we presented users with a list of one hundred science stories drawn from the Science section of the New York Times and science-related press releases from major universities (sciencenewsdaily.com). We asked each user to recommend articles they found interesting. Each article contains a title, short description, and a link to the full story. We recorded all user actions, including votes and URL clicks.

The design was similar to that of the landmark experiments conducted by Salganik et al. [Salganik et al. 2006], which presented users with a set of songs arranged either in a random order or ordered by their popularity. In the latter case, the popularity was shown to users, thereby providing a social signal. This led to large and unpredictable variations in song popularity, which they attributed to social influence. Other experiments demonstrated that social influence biases decisions, producing "irrational herding" and "rich get richer" effects [Salganik and Watts 2008; Lorenz et al. 2011; Muchnik et al. 2013].

We tested four policies for ranking stories shown to users. The *random* ranking presented the stories in a random order, with a new ordering generated for each user. The *popularity* ranking ordered stories by their popularity, i.e., in the decreasing order of the number of votes, and the *activity* ranking ordered them in chronological order of the latest votes, with the most recently recommended story at the top of the list. Finally, the *fixed* ranking showed all stories in the same order to every user.In the *social influence* condition, we also showed users the number of votes each story had previously received.

The random ranking serves as a control by averaging over position, so an item's popularity reflects its quality [Salganik et al. 2006]. To characterize differences between items, we define the *appeal* a_s of a story s to a population of users as the conditional probability a user who sees story s subsequently votes for it. The appeal of each story allows estimating the number of votes we would expect at each position in the random ranking. Comparing this value to the actual number of votes stories received gives the position bias, Figure 1(a). Users pay five times more attention to top stories than those in the middle of the list.

Controlling position allows directing users' attention, which has a dramatic effect on the evolution of story popularity. Figure 1(b) shows the correlation between story popularity and appeal as a function of the number of recommending users. The activity ranking is better (and faster) at uncovering quality stories than the popularity ranking, resulting in higher correlation with appeal. Thus, to highlight interesting content to users, content providers should rank items by recency of recommendation, rather than the common choice to rank by popularity.

Inequality of outcomes, i.e., popularity, is measured by the Gini coefficient (Fig. 1(c)). In the random ranking, inequality arises from the variation in story quality. Activity ranking produces slightly more inequality than expected from quality variations. On the other hand, the fixed and popularity rankings produce the most inequality (Gini of popularity ranking is 0.409). Unlike conclusions of Salganik et al., large inequality can exist in the absence of social influence, when popularity ranking leads people to pay more attention to popular stories, producing a "rich get richer" effect.

Social influence further increases the inequality of outcomes by directing attention to popular stories and also increasing user preference for popular stories. Figure 1(d) compares observed votes to the expected numbers of votes those stories would receive when there is no social influence. Stories associated with small influence signals (bottom quartile) get fewer votes than expected, and those with signals in the top quartile get about twice as many votes as those in the bottom quartile. Thus, social influence is half as important as position bias in explaining variation in popularity. Social influence

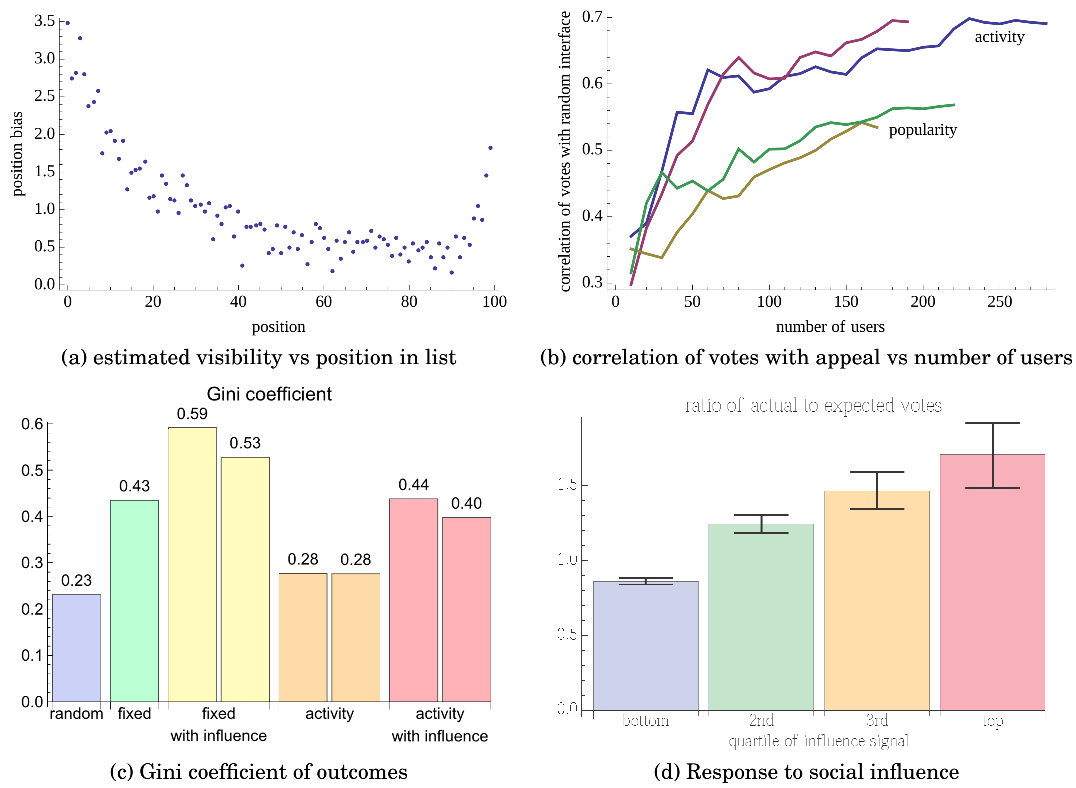

(a) estimated visibility vs position in list

(b) correlation of votes with appeal vs number of users

(c) Gini coefficient of outcomes

(d) Response to social influence

Fig. 1. Experimental results.

changes the effort users devote to recommendation. With influence, users spent on average about 20% less time on the task than without the influence signal. However, this does not appear to compromise recommendation quality, resulting in similar correlation of popularity with quality.

3. DISCUSSION

Our experiments show that accounting for human cognitive biases help us better understand — and control — collective behavior. Attention, which is mediated by cognitive biases, plays an important role in collective dynamics of peer recommendation, with most of variation coming from position bias. Knowing how much attention an item receives allows for better estimates of its underlying quality from the observed votes. This, in turn, enables predicting the growth of items' popularities, i.e., whether they will become blockbusters or fail to gain much user interest. More interestingly, rather than merely predict item popularity from the response of early users, we can aim to control this behavior. Specifically, we showed that by manipulating the attention items receive, merely by changing their position in a list, we can improve the outcomes of peer recommendation and more efficiently identify quality items.

Integrating social influence into group decision making processes can lead to trade-offs in performance. While influence makes it easier for a group to achieve consensus and adopt new ideas, the bias it introduces into individual judgements may skew the outcomes of collective computation tasks that rely on independent decisions of many people. However, our work shows the benefits of social influence: by reducing effort, social influence makes it easier to collect opinions of more people, without significantly sacrificing performance of peer recommendation.

REFERENCES

Niels J. Blunch. 1984. Position Bias in Multiple-Choice Questions. *Journal of Marketing Research* 21, 2 (1984), pp. 216–220.

Georg Buscher, Edward Cutrell, and Meredith R. Morris. 2009. What do you see when you're surfing?: using eye tracking to predict salient regions of web pages. In *Proc. the 27th Int. Conf. on Human factors in computing systems*. 21–30.

Kristina Lerman, and Tad Hogg. 2014. Leveraging Position Bias to Improve Peer Recommendation. *PLoS One* 9(6), (2014), e98914.

J Lorenz, H Rauhut, F Schweitzer, and Dirk Helbing. 2011. How social influence can undermine the wisdom of crowd effect. *Proceedings of the National Academy of Sciences* 108, 22 (2011), 9020–9025.

Lev Muchnik, Sinan Aral, and Sean J. Taylor. 2013. Social Influence Bias: A Randomized Experiment. *Science* 341, 6146 (2013), 647–651.

Stanley L. Payne. 1951. *The Art of Asking Questions*. Princeton University Press.

Matthew J. Salganik, Peter S. Dodds, and Duncan J. Watts. 2006. Experimental Study of Inequality and Unpredictability in an Artificial Cultural Market. *Science* 311, 5762 (2006), 854–856.

Matthew J. Salganik and Duncan J. Watts. 2008. Leading the Herd Astray: An Experimental Study of Self-fulfilling Prophecies in an Artificial Cultural Market. *Social Psychology Quarterly* 71, 4 (1 Dec. 2008), 338–355.