

Social Computing and User-generated Content: A Game-theoretic Approach

ARPITA GHOSH
Cornell University

Social computing is now ubiquitous on the Web, with user contributions on sites like online review forums, question-answer forums, wikis, or Youtube forming a growing fraction of the content consumed by Web users. But while there is plenty of user-generated content online, the quality of contributions and extent of participation vary widely across sites. We survey some recent work taking a game-theoretic approach to the problem of incentivizing high quality and participation in these social computing systems.

Categories and Subject Descriptors: J.4 [**Social and Behavioral Sciences**]: Economics

General Terms: Social Computing, Crowdsourcing, User-generated Content (UGC), Game theory

1. INTRODUCTION

The Web is increasingly centered around contributions by its users. User-generated content (UGC) such as Amazon and Yelp reviews, Wikipedia articles, blogs, or YouTube videos now constitute a large fraction of the relevant, easily accessible content that makes the Web useful, and crowdsourcing tasks to the online public is increasingly common, ranging from systems based on unpaid contributions such as Games with a Purpose or online Q&A forums (Y! Answers, Quora, and StackOverflow to name a few), to platforms for paid crowdsourcing such as Amazon's Mechanical Turk and TopCoder. But while some websites consistently attract high-quality contributions, other seemingly similar sites are overwhelmed by junk, and yet others fail due to too little participation. There is a growing body of work in the social psychology literature on what factors motivate, or constitute rewards for, participants in these social computing systems, and on user experience and interface design to exploit these factors. But these rewards are nonetheless not unlimited resources, and must be distributed appropriately amongst participants to incentivize desirable behaviors. Given the understanding from the social psychology literature on what constitutes a reward for these contributors, how can we design the *allocation* of these rewards to incentivize desirable outcomes?

A key aspect to modeling, analyzing, and finally designing mechanisms for these social computing systems is to recognize that participation in all these systems is *voluntary*— contributors have a choice whether to participate in the system at all and indeed, many UGC sites fail, either immediately or eventually, from too few contributions. Second, even after having decided to participate, contributors can decide how much effort to put into their contributions, which affects the quality of the output they produce. Call a method to allocate these rewards — either monetary as in paid crowdsourcing, or non-monetary such as attention [Huberman

Email address: arpitaghosh@cornell.edu

et al. 2009], reputation and status [Beenen et al. 2004], or virtual points (which, quite miraculously, do appear to motivate contributors [Nam et al. 2009; Yang and Wei 2009; Yang et al. 2011])— amongst contributors an *incentive scheme*. What can we understand, using a game-theoretic approach, about what incentive schemes a website should employ to consistently elicit high quality contributions from self-interested agents?

2. ATTENTION REWARDS

The first ingredient essential to addressing such problems is a *model* within which the problem of analyzing and designing mechanisms for social computing systems can be formalized. In [Ghosh and McAfee 2011], we take the first steps towards such a model in the context of diverging *attention* rewards, as in UGC sites with high viewership. The model consists of strategic contributors motivated primarily by exposure, or viewer attention¹, with the feature that the quality *as well as* the number of contributions are endogenously determined in a free-entry Nash equilibrium— we emphasize this endogenous choice to contribute since the problem in UGC is not just limited to incentivizing high quality but also *encouraging* the production of content, and a model with exogenous participation misses a salient factor in most UGC settings.

Without some connection between quality and exposure, exposure-motivated contributors will flood a site with low quality contributions, as is indeed observed in practice. Is there a way to allocate the available attention from viewers (by choosing which contributions to display on the webpage and with what prominence) amongst the contributions— a *mechanism*— that encourages high quality, while also maintaining a high level of participation, in equilibrium? An ideal mechanism in this context would be easily implementable in practice and elicit a large number of high quality contributions, with near-optimal quality as the available attention diverges: the diverging attention regime corresponds to the large viewership typically associated with UGC, and it is these high-traffic sites that most attract such attention-motivated contributors. We demonstrate that a very simple elimination mechanism can indeed achieve quality that tends to optimal, along with diverging participation, as the number of viewers diverges.

But what about the mechanisms that are actually used by online UGC sites? Many websites attempt to rank content according to ‘quality’ estimates from the thumbs-up/thumbs-down style ratings by viewers, displaying higher-rated contributions more prominently by placing them near the top of the page and pushing less well-rated ones to the bottom. The model in [Ghosh and McAfee 2011] turns out to be very useful to analyze equilibrium behavior in this widely-used rank-order mechanism as well. In [Ghosh and Hummel 2011], we show that this rank-order mechanism also elicits high quality contributions, in a very strong sense, while still achieving high participation— the *lowest* quality that can arise in *any* mixed strategy equilibrium of the rank-order mechanism becomes optimal as the amount of

¹Note that this does not rely on any assumption about *why* contributors seek greater exposure—it could be because of a type of altruism (a contributor believes her contribution is valuable, and wants it to benefit the largest number), or a contributor seeking the largest possible audience for self-expression, or a host of other social-psychological factors.

available attention diverges. Additionally, these equilibrium qualities are higher (with probability tending to 1 as the amount of available attention diverges) than those elicited in equilibria of a more equitable (and less draconian) proportional mechanism, which distributes attention in proportion to the number of positive ratings a contribution receives.

3. CROWDSOURCING: CONTENT AND CONTESTS

A somewhat different incentive design problem crops up in crowdsourcing settings with finite rewards. Crowdsourcing, where a problem or task is broadcast to a ‘crowd’ of potential contributors for solution, is seeing rapid and wideranging adoption online, all the way from companies seeking solutions to their projects using crowdsourcing contests run by Innocentive or TopCoder, to crowdsourced *content* solicited by individuals asking questions on online Q&A forums like Y! Answers, StackOverflow or Quora. In these crowdsourcing settings, there is typically some finite prize to be distributed amongst participants, unlike the diverging reward regimes that arise in the context of attention rewards on UGC sites with huge viewership. While there is a growing literature on the optimal design of contests as well as online crowdsourcing contests (see §1 in [Ghosh and McAfee 2012] for a survey), this literature assumes some *fixed* number of contestants who always participate. However, because there is a nonzero cost to making a contribution of *any* quality which can be avoided by simply not participating at all (and indeed many sites based on crowdsourced content do not have adequate participation), it is more realistic to model entry as a strategic choice. The final number of contestants is then not a fixed number given apriori, but rather endogenously determined in equilibrium. How should rewards be designed to incentivize high effort in crowdsourcing settings when entry is an *endogenous*, strategic choice?

In [Ghosh and McAfee 2012], we investigate designing incentives in environments with endogenous entry for two kinds of rewards that arise in the context of crowdsourcing. The first is in the context of bounded attention rewards in online Q&A forums like Quora or StackOverflow, where the mechanism designer, or site owner, has a choice about which of the received answers to display for each question, and how— he could choose to display all answers for a particular question or display only the best few and suppress some of the poorer contributions (by either not displaying them at all, or by ‘fractionally’ displaying them some of the time, for example, rotating amongst these answers). On the one hand, suppression should cause quality to rise, because the payoff to poor content falls; on the other hand, suppressing content also corresponds to decreasing the total reward paid out, which could decrease quality. What strategy improves the quality of the best contribution supplied, and what about average quality? It turns out that here, the *entire* equilibrium distribution, and therefore every increasing statistic including the maximum and average quality (accounting for participation), improves when the rewards for every rank but the last are as high as possible— specifically, when the cost of producing the lowest possible quality content is low, the optimal mechanism displays all but the poorest contribution.

A different constraint arises in crowdsourcing contests with monetary rewards, where the principal posing the task offers some fixed amount of money for the

solution— here, there is some fixed total available reward which can be distributed *arbitrarily* amongst the agents. (The distinction between this setting and the previous one is that it is not possible to take away attention from a lower position and ‘add’ it to a higher position since, to a first approximation, attention to lower spots comes from a subset of viewers providing attention to the higher ones; so attention rewards could not be arbitrarily redistributed across ranks.) When entry is exogenous, quality increases with the number of participants available, suggesting that subsidizing entry to increase participation may be productive in the endogenous entry setting as well. And even if subsidizing entry (at the cost of paying less to the winner) were to ‘reduce’ the equilibrium distribution from which each contributor chooses her quality, the *expected* value of the maximum quality might nonetheless increase when the number of contributors increases, since we have the maximum of a larger number of random variables. How does this tradeoff work out? It turns out that subsidizing entry actually does not improve the expected value of the maximum quality, although it may improve the average quality— it turns out that when entry is endogenous, free entry (corresponding to a winner-take-all contest) is dominated by *taxing* entry, *i.e.*, making all entrants pay a small fee which is rebated to the winner, along with whatever rewards were already assigned, can improve the expected quality of the best solution.

Implementability. In [Ghosh and McAfee 2012] we ask what is the best outcome, given certain resource constraints, that can be obtained in equilibrium when participants are strategic agents. A different question is whether the optimal outcome that can be achieved with *non-strategic* agents (for some objective of the mechanism designer) can be at all implemented in an equilibrium of *any* mechanism in the presence of strategic agents. In [Ghosh and Hummel 2012], we investigate implementability in the context of social computing sites like online Q&A forums. A number of these forums (Y! Answers, MSN QnA, and Rediff Q&A, to name a few) are structured so that the ‘best’ contribution for each task receives some high reward, while all remaining contributions receive an identical, lower reward irrespective of their actual qualities. Suppose a mechanism designer (site owner) wishes to optimize an objective that is some function of the number and qualities of received contributions. When potential contributors are strategic agents, who decide whether to contribute or not to selfishly maximize their own utilities, are optimal outcomes at all implementable, and if yes, is such a ‘best contribution’ mechanism adequate to implement an outcome that is optimal for the mechanism designer? We find that when a contribution’s value is determined primarily by contributor expertise (as in forums for medical, legal or other questions requiring domain expertise) and agents only strategically choose whether or not to contribute, such best contribution mechanisms, quite surprisingly, *can* implement optimal outcomes— for any reasonable objective of the mechanism designer, the relative rewards for the best and remaining contributions can always be chosen so that the outcome in the unique symmetric equilibrium of the mechanism maximizes this objective. However, the situation is somewhat more subtle when a contribution’s value depends on both the contributor’s expertise and the effort she puts in, so that both participation and effort are endogenous: if the system can rank the qualities of contributions perfectly, optimal outcomes can never be implemented by contests.

However, with adequate *noise* in the contributions' rankings, agents will again use strategies that maximize the mechanism designer's utility in equilibrium— thus imperfect rankings can actually help achieve implementability of optimal outcomes when effort is endogenous and influences quality, a result somewhat reminiscent of the use of randomization in circumventing impossibility results in algorithm design.

4. FURTHER DIRECTIONS

There are a number of exciting directions for further work to develop the game-theoretic foundations for social computing systems, both extending our models to relax simplifying assumptions, as well as mechanism and market design questions requiring entirely new models.

First, the model of content we use is fairly simple. A contribution's quality is modeled as one-dimensional; also, the value derived from a *set* of contributions, rather than a single contribution, is not modeled at all, nor is the possibility that different users derive value from different contributions. Arguably, a multi-dimensional model of quality is a more realistic representation of the value of a single contribution (although such a multi-dimensional model will bring up the problem of modeling users' one-dimensional ratings of contributions, as is typical on many websites, in terms of these multiple dimensions). A model with such vector qualities could also help define the value from a set of contributions, as well as capture differing user preferences over contributions. Finally, viewer ratings of content, assumed thus far to be *probabilistically* dependent on quality, might perhaps more accurately be modeled in some settings as the outcome of a *strategic* choice, since contributions are often rated by competing contributors. Modeling and designing mechanisms for systems where both the contributors and raters of content are strategic agents is a promising direction for further research.

A second dimension that is unexplored by the current literature relates to the temporal aspect. At the level of individual tasks, we assume that agents make simultaneous choices about their contributions, whereas many UGC environments are perhaps better suited to an alternative (although harder to analyze) sequential model in which potential contributors arrive at different times and make decisions about their own contributions after viewing the existing set of contributions. What mechanisms elicit adequate participation and quality in such dynamic sequential models? A related family of problems at the interface of incentives and learning regards mechanisms for learning the qualities of contributions from viewer feedback in the presence of strategic agents, both for a single task (a first formulation for this is presented in [Ghosh and Hummel 2013]), as well as designing user reputations using performance in multiple tasks to incentivize consistent and frequent high-quality contributions. Moving beyond individual tasks and users, an interesting direction relates to investigating an entire *site's* evolution over time in terms of attracting and sustaining adequate contributions and contributors, with models that both explain observed site growth and decline as well as allow *designing* mechanisms for sustained participation and quality over time. Finally, a central question in this line of work is developing a more nuanced understanding of contributors' cost and benefit functions, using experimental and empirical studies to incorporate increasingly accurate models of contributor motivations into a game-theoretic framework.

REFERENCES

- BEENEN, G., LING, K., WANG, X., CHANG, K., FRANKOWSKI, D., RESNICK, P., AND KRAUT, R. 2004. Using social psychology to motivate contributions to online communities. In *Proceedings of the 2004 ACM conference on Computer supported cooperative work (CSCW'04)*.
- GHOSH, A. AND HUMMEL, P. 2011. A game-theoretic analysis of rank-order mechanisms for user-generated content. In *12th ACM conference on Electronic commerce (EC'12)*.
- GHOSH, A. AND HUMMEL, P. 2012. Implementing optimal outcomes in social computing. In *21st International World Wide Web Conference (WWW'12)*.
- GHOSH, A. AND HUMMEL, P. 2013. Learning and incentives in user-generated content: Multi-armed bandits with endogenous arms. In *4th Conference on Innovations in Theoretical Computer Science (ITCS'13)*.
- GHOSH, A. AND MCAFEE, R. 2011. Incentivizing high-quality user generated content. In *20th International World Wide Web Conference (WWW'11)*.
- GHOSH, A. AND MCAFEE, R. 2012. Crowdsourcing with endogenous entry. In *21st International World Wide Web Conference (WWW'12)*.
- HUBERMAN, B., ROMERO, D., AND WU, F. 2009. Crowdsourcing, attention and productivity. *Journal of Information Science*, 35(6):758-765.
- NAM, K., ACKERMAN, M., AND ADAMIC, L. 2009. Questions in, knowledge in? a study of naver's question answering community. *ACM Conference on Human Factors in Computing Systems (CHI'09)*.
- YANG, J., ACKERMAN, M., AND ADAMIC, L. 2011. Virtual gifts and guanxi: Supporting social exchange in a chinese online community. *ACM Conference on Computer Supported Cooperative Work (CSCW'11)*.
- YANG, J. AND WEI, X. 2009. Seeking and offering expertise across categories: A sustainable mechanism works for baidu knows. *International AAAI Conference on Weblogs and Social Media (ICWSM'09)*.