CourseRank: A Closed-Community Social System Through the Magnifying Glass

Georgia Koutrika, Benjamin Bercovitz, Filip Kaliszan, Henry Liou, Hector Garcia-Molina
Computer Science Department, Stanford University, California, USA
{koutrika, berco, filip.kaliszan, liouh, hector}@cs.stanford.edu

Abstract

Social sites are extremely popular among users but user interactions in most sites revolve around relatively simple tasks, such as uploading resources, tagging and poking friends. We believe that social sites can go beyond simple interactions among individuals and offer valuable services to well-defined, closed, communities (e.g., an academic, corporate or scientific community). In this paper, we present an example of a *closed-community social system*, CourseRank, an educational and social site where Stanford students can explore course offerings and plan their academic program. We perform an analysis of 12 months worth of CourseRank data including user contributed information, such as ratings and comments, as well as information extracted from the user logs, and we analyze several aspects of user interactions and user-contributed content in the site, such as activity levels, user behavior and user content quality. Our findings provide useful insights with respect to the potential of closed-community social sites.

Introduction

Social web sites, such as FaceBook, del.icio.us, Y! Answers, Flickr and MySpace, have become important components of the Web. In these sites, a community of users contribute resources, which can be photos, personal information, evaluations, votes, answers to questions or annotations. Social sites have become extremely popular among users because "people start to understand that they can publish virtually anything and put it on the web for anyone to see if so they wish, and they are their own broadcasters" (Coelho 2008).

Currently, most social sites mainly focus on resource sharing among web users and user interactions revolve around relatively simple tasks, such as uploading resources (e.g., photos and videos), tagging (i.e., adding simple, descriptive words to existing resources), connecting to other people, and so forth. We believe that social sites can go beyond simple interactions among individuals and offer valuable services to well-defined, closed, communities (e.g., an academic, corporate or scientific community). These specialized social sites may have official data (e.g., corporate documents and forms, scientific papers, course bulletins, etc) in combination with user contributed information providing an added-value communication and interaction environment.

In this paper, we present such a *closed-community social system*, CourseRank. This is an educational and social site where Stanford students can explore course offerings and plan their academic program. Students can explore official information provided by the university as well as information provided by students, such as comments and ratings for courses. Faculty members and university administrators can also participate, providing useful information for students. We perform an analysis of CourseRank data we have collected from the first year of its deployment, including user contributed information, such as ratings and comments, as well as information extracted from the user logs. We analyze several aspects of user interactions and user-contributed content in the site, such as activity levels, user behavior, user content quality, and so forth. Our analysis aims at examining some important hypotheses and facts that widely hold in general-purpose web (social) sites. Our findings provide interesting insights in how the combination of non-social information with social features can lead to a well-adopted, successful, community site.

Outline. The paper is organized into the following sections: (a) a review of related work on social sites; (b) an overview of CourseRank; (c) a study of the usage patterns of the system; (d) an analysis of user activity and behavior; and (e) a discussion of ingredients for a successful special-purpose social site.

Related Work

While most social sites are open to the world wide web and are of great benefit for publicly accessible resources, such as photos (e.g., Flickr), URLs (e.g., Del.icio.us) and research papers (e.g., CiteULike), there are social sites that target closed communities with possibly restricted resources, such as enterprise social sites (Millen, Feinberg, and Kerr 2006). In this paper, we show another example of a special-purpose social site for the university community.

Understanding user behavior is a crucial step in building more effective systems and this fact has motivated a large amount of user-centered research on different web-based systems (Adar et al. 2007; Golder and Huberman 2006; Marlow et al. 2006; Sen et al. 2006; Xu et al. 2006; White and Drucker 2007). A number of studies focus on

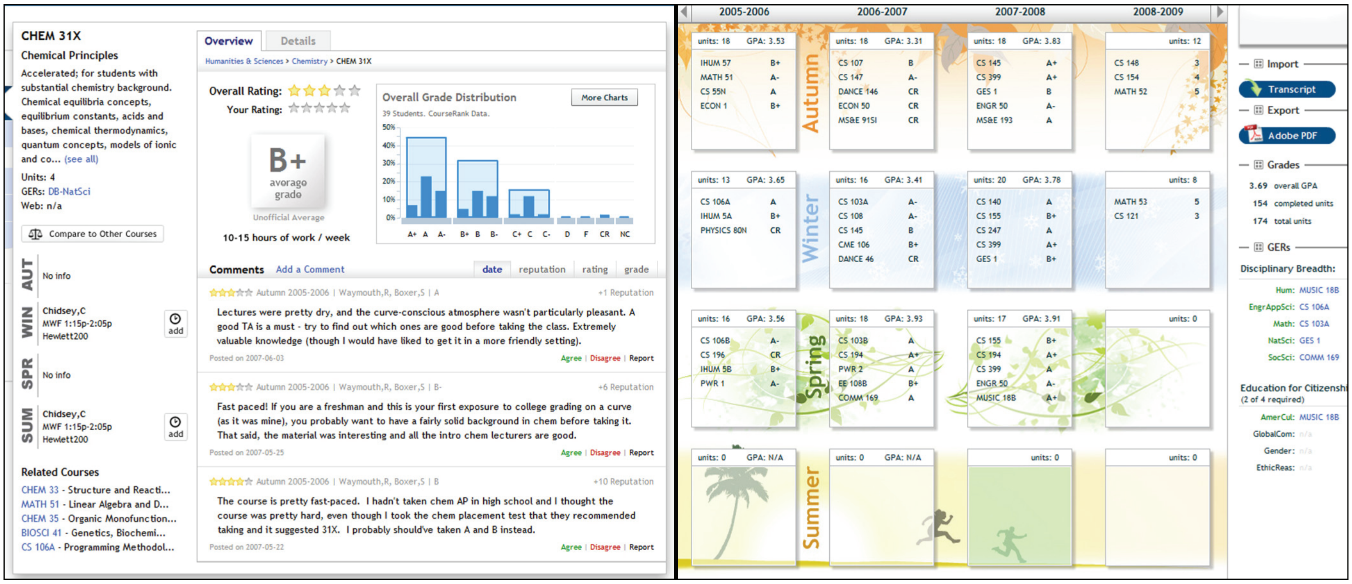

Figure 1: CourseRank Screen Shots: course description (left), course planner (right).

user tagging behavior in social systems. For instance, an experimental study of tag usage in My Web 2.0 shows that people naturally select some popular and generic tags to label Web objects of interest (Xu et al. 2006), while other studies identify factors that influence personal tagging behavior, such as people's personal tendency to apply tags based on their past tagging behaviors and community influence of the tagging behavior of other members (Golder and Huberman 2006; Marlow et al. 2006; Sen et al. 2006). Other studies have been performed on question-answers sites (Adamic et al. 2008; Gyöngyi et al. 2008) and social networking sites (Lampe, Ellison, and Steinfield 2008). Our study reveals many unexpected patterns and trends that do not generally hold in other social sites and are shaped to a great extent by the fact that the examined site targets the special needs of a closed community.

CourseRank

CourseRank is a social site where Stanford students can review courses and plan their academic program by accessing official university information and statistics, such as bulletin course descriptions and grade distributions. Students can also provide information, such as comments on courses, ratings, questions and answers. To illustrate, the system provides (January 2009) access to $18,605$ courses, $134,000$ official evaluations, and over $50,300$ ratings. The system is already used by approximately 10,000 Stanford students out of a total of about 14,000 students.

Using CourseRank, students can search for courses of interest, rank the accuracy of each others' comments and get personalized recommendations. They can shop for classes, and organize their classes into a quarterly schedule or devise a four year plan. CourseRank also functions as a feedback tool for faculty and administrators, ensuring that information is as accurate as possible. Faculty can also modify or

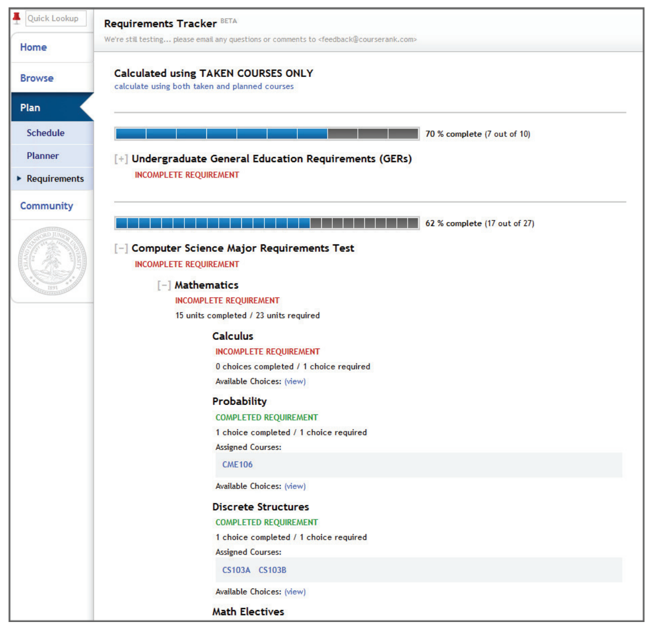

Figure 2: Requirements Tracker

add comments to their own courses, and can see how their class compares to other classes. Figure 1 shows two CourseRank screen shots: on the left is part of a course descriptor page, and on the right is the 4-year course planner[1].

CourseRank has several important features that distinguish it from classical social sites but also from other public course evaluation sites (e.g., `ratemyprofessors.com`).

- It provides access to both official Stanford data (e.g., course descriptions, schedules and results of course eval-

[1] At our site, (`http://courserank.com`), visitors can see a video with student testimonials and a demo (demo tab).

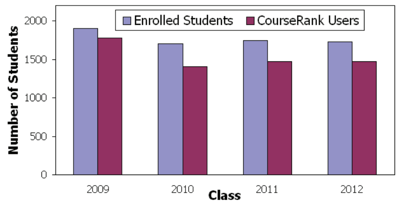

Figure 3: User population (undergrads)

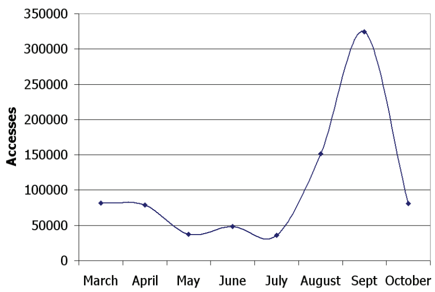

Figure 4: Usage

uations conducted by the university), as in a typical database application, as well as to user-contributed information (e.g., course rankings, comments and questions), as in a typical social system.

- It provides many tools similar to ones found at existing social sites for searching, evaluating courses and other users' opinions, getting recommendations from the system and asking questions and providing answers. In addition, CourseRank offers powerful tools geared to our domain. For example, students can check if the courses they have taken (or are planning to take) satisfy the requirements for their major. Figure 2 provides a screen shot from the requirement tracker page.
- It is a closed-community social site since it is only available to the Stanford community. It has access to official "user names" on the Stanford network and can therefore validate that a user is a student or a professor or staff. In addition, unlikely the "open" Web, all data is centrally stored and we have control over the site. There are other examples of sites with similar characteristics. For example, a corporate social site, where employees and customers can interact and share experiences and resources, shares many features with CourseRank: the need to service a varied constituency (employees, managers, customers, etc), restricted access, and so forth.

Service Evolution

CourseRank has been released in 2007 and it has undergone two major updates, one early 2008 and the other in September 2008. How popular has CourseRank been among students? What is the impact of its various tools? What usage patterns we observe in CourseRank?

Popularity. Figure 3 shows how many Stanford undergraduate students enrolled in the classes of 2009 to 2012, i.e., who are expected to graduate in these years, have signed

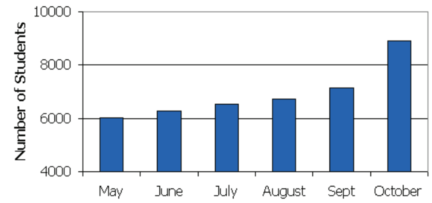

Figure 5: Users signing up with the system

up with our course planning site. Stanford accepts around 1700 freshmen every year for the last few years. We see that around 85% of each class are registered in CourseRank.

An important factor for CourseRank's success has been the fact that Stanford University is providing us with useful data for the site. Students can review course descriptions, schedules and results of official course evaluations conducted by the university. The combination of official and user-contributed data adds value to the system. Furthermore, the Registrar has assigned a person to be our CourseRank liaison and to make sure that we get data on a timely basis. This commitment shows that not only students but the administration also views CourseRank as an important tool that needs to be supported long term.

Usage. Figure 4 shows the usage of CourseRank for the period from March 2008 to Oct 2008 (the site's current structure has been stabilized right before March 2008). We observe an interesting pattern: CourseRank usage follows well user needs during an academic year. Hence, we see high traffic at the beginning of each quarter, i.e., March and September for the period we study, where students enroll to classes, and the lowest traffic between May and July, i.e., when the academic year ends. High traffic is observed close to the beginning of the academic year starting in August and culminating in September due to two facts. First, students try to find which classes to take and make a plan not just for the quarter but possibly for the whole academic year. Second, freshmen arrive and they need orientation and familiarizing with the available learning options in the university.

Figure 5 shows the number of active users per month and allows to take a closer look at user registrations in the system. We see that the number of users in October has considerably increased compared to the numbers in September and August, which is due to the arrival of freshmen. It is surprising that many new students sign up with the system even before arriving at Stanford (starting from August). This fact serves as an indication of the usefulness and impact of CourseRank as a focused social site. It may be even more interesting to also observe that users sign up with the system all year round not only at the beginning of a new quarter or of the academic year. This constant flow of new users entering the system provides more evidence of its usefulness. It may also indicate a community effect: based on user testimonials, students that hear about CourseRank from other students decide to join.

Tools Popularity. CourseRank offers various tools. For example, the course pages (*Course*) provide useful information for a course, such as the course description, grade dis-

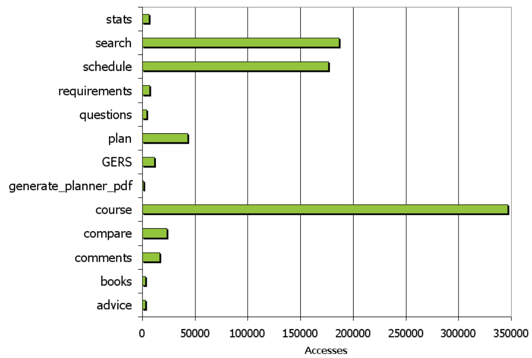

Figure 6: Popular tools

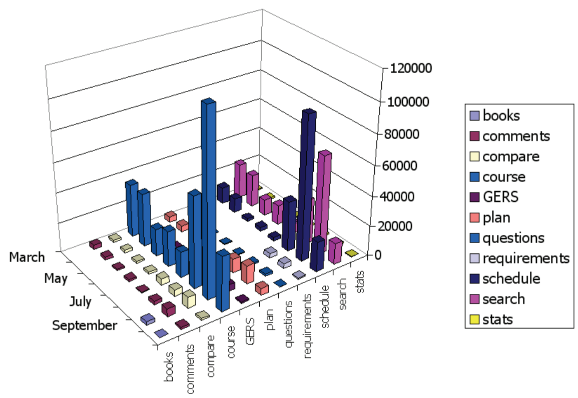

Figure 7: Popular tool usage

tributions, user ratings, comments, and so forth. Figure 6 shows the popularity of these tools for the same period (from March 2008 to Oct 2008). We observe that the course pages (*Course*) are the most popular destination in the site, followed by the search feature (*Search*), the weekly schedule (*Schedule*) and the planner (*Plan*). Comparing courses (*Compare*) and writing comments for courses (*Comments*) have their share of traffic.

Figure 7 provides a detailed view of the tools' usage over the same period of time. We observe that the tools for checking requirements (*GERS*), exchanging books (*Books*), and the Question and Answer forum (*Questions*) have been added later in the system. Interestingly, not all features have been popular to date. The Question and Answer forum has little traffic because there are no incentives to visit: If there are few questions or answers, why would people ask questions or go looking for answers there? To address this shortcoming, we plan to seed the forum with "frequently asked questions" developed in conjunction with department managers, e.g., "who do I see to have my program approved?" or "what is a good introductory class in department X for non-majors?" Questions will be automatically routed to people who are likely to be able to answer them. With a useful body of questions and answers, we hope students will start using the forum.

Analysis

We analyze the activity, behavior and profiles of users in the system in order to gain insights into the impact and the merits of the system. Our analysis evolves around a number of hypotheses based on what is widely known in general-purpose web (social) sites as well as based on what one might expect to find in a closed-community social site. In particular, we are testing the following hypotheses:

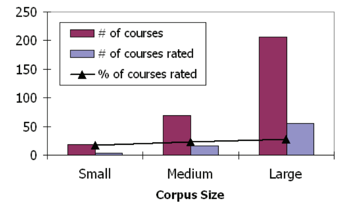

Figure 8: User ratings vs. corpus size

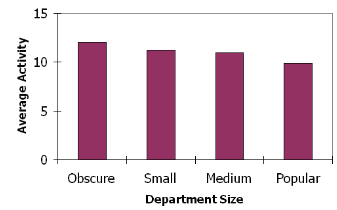

Figure 9: User activity vs. department size

H1: Small (collection) is better
H2: The 90-9-1 rule rules users
H3: The community effect
H4: Power-law distributions everywhere
H5: Social sites are for techies
H6: Diverse user interests
H7: The objectivity principle
H8: Spam is everywhere
H9: Everyone lies

User Activity

H1: Small (collection) is better? In social sites with large collections, it may be hard to find and rate items of interest. Are users discouraged by large corpus sizes? Are smaller collections more manageable and hence attract user input?

We group the university's departments based on the size of their course corpus to: *small* departments (offering less than 30 courses), *medium* departments (offering between 30 and 100 courses) and *large* departments (offering between 100 and 200 courses). Figure 8 shows the average number of courses offered and the average number of courses rated for each group. The black line shows the percentage of courses rated within each group. Interestingly, the number of visible and popular courses (in terms of attracting user input not high ratings) does not shrink with the corpus size. Larger departments offer more courses but they also have more students. The collective contributions of a larger number of users covers a larger part of the corpus.

H2: The 90-9-1 rule rules users? In 2006, Jakob Nielsen coined the phrase *participation inequality*, referring to participation by online communities. He says that online social site activity is generated largely by a small number of the

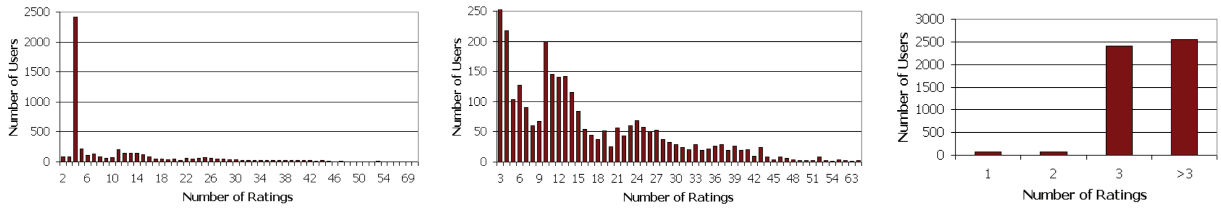

(a) distribution of ratings (b) the long tail (c) the power of the long tail

Figure 10: User activity

community (Nielsen 2006). Nielsen sums up participation inequality in the 90-9-1 rule: 90% of registered users do not contribute (often called lurkers), 9% of users contribute from time to time and 1% of users participate a lot and account for most contributions. For example, in Y! Answers, only a small fraction (< 18%) of the users provides feedback in the system in the form of votes (Gyöngyi et al. 2008). In Y! Groups, 90% of the users registered are lurkers (Horowitz 2006).

We now group departments based on the number of their students to: *obscure* with less than 20 students, *small* with a number of students between 20 and 100, *medium* with over 100 students and less than 300 and *popular* with over 300 students. We measured the average percentage of active users per program type, and we found it to be around 65.13% in all cases (with only very small variations). This percentage for active users is remarkably higher compared to other general-purpose social sites.

CourseRank is a social site for a small closed community offering tools geared to our domain. We believe that the value of these tools in helping students organize their academic program makes users more actively engaged. Furthermore, the existing structure in the university (or other organizations for that matter) - unlike online communities on the open web - helps influence people to become more active and willing contributors. In our case, students encourage other students to use CourseRank. We drill more down to the community effect below.

H3. The community effect? Do the users of larger departments contribute more? Does the size of the community affect user activity?

We again group departments based on the number of their students as above. Figure 9 shows that the average number of ratings a user contributes decreases as the size of the group the user belongs to increases. We think that this phenomenon is due to the effect of the community size. Smaller communities breed more active users because the feeling of belonging to a community is stronger. Affinity is the key driver in forming online communities. In case of a university or an organization, it may also be the case that people of small communities may know each other in real life too.

H4. Power-law distributions everywhere? Many systems and phenomena are distributed according to a power law distribution. A power law applies to a system when large is rare and small is common.

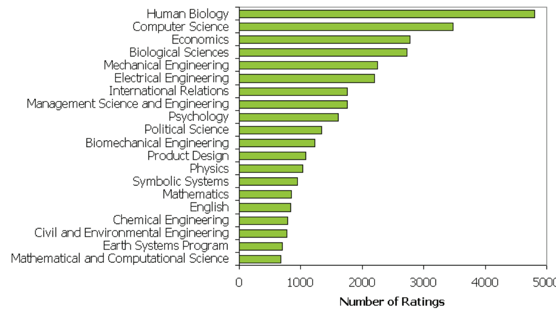

Figure 11: Number of ratings

Figure 10(a) shows the distribution of ratings per user in the system. We observe that a large number of users have given only 3 ratings. This is due to the fact that CourseRank asks undergrads (except freshmen, who are therefore not shown at all in this figure) to rate three courses they have previously taken. If we zoom in on the long tail of the distribution (Figure 10(b)), we observe that the user contribution does not decrease smoothly but there are peaks. Remarkably, this distribution has a wavy form and does not resemble the power-law distributions observed in typical social sites. A reason for that is that many students enter groups of course ratings at the beginning of each quarter.

Finally, Figure 10(c) shows that despite the fact that a large number of users contributed only 3 ratings, there are more users that contribute more ratings in the system. Consequently, the power of the long tail is considerable.

User Profiles

H5: Social sites are for techies? In a focused social site, such as CourseRank, users have well-defined interests and background. How do the user profiles shape user behavior? What is the profile of active participants in CourseRank? Are people interested in computer-related programs more inclined to get actively involved?

We group users based on the academic program they follow. Since freshmen in Stanford University have not signed up to a particular program yet, this analysis shows patterns in the remaining student base (i.e., sophomores, third-year students, and so forth). Furthermore, we consider only programs with over 30 students. Programs with few students provide a very small sample that does not help obtain a realistic view of user activity levels. For example, a program with 5 active users out of 5 students, will appear as a pro-

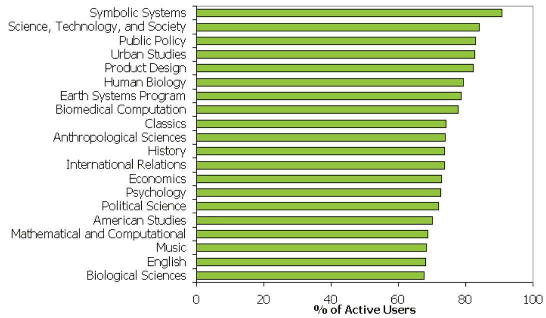

Figure 12: Active users

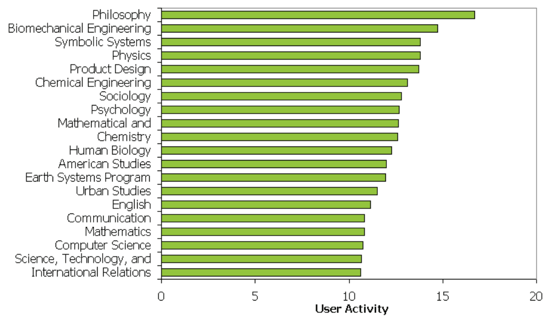

Figure 13: Average user activity

gram with many active users, as a program with 200 active users out of 200 students. Figures 11, 12 and 13 provide different views of user activity based on the academic program.

Figure 11 shows the total number of ratings provided by students in the top 20 programs in terms of aggregate user activity. The figure seems to confirm a lurking suspicion about "techies mostly participating in social sites". Indeed, computer-oriented and engineering programs, such as Computer Science, Electrical Engineering, and Management Science and Engineering, dominate the top 20 programs in terms of aggregate user activity. However, looking deeper into user activity levels reveals a different truth.

In particular, Figure 12 shows the percentage of active users within each academic program. The figure shows only the top 20 programs based on the number of active users and unveils an unexpected surprise. The typical computer-oriented programs (e.g., Computer Science and Electrical Engineering) are not within the top 20 programs. On the other hand, more classical fields, such as Urban Studies, Public Policy, History, Psychology and Classics, appear to have a greater percentage of active students.

Furthermore, Figure 13 shows the average number of ratings per individual for each program for the top 20 programs ordered on individual user activity. We observe that the most active users are not necessarily involved with computers at an academic level but are interested in studies on other areas, such as Philosophy, Sociology and American Studies. In fact, computer science students are in the tail of the most active users. This fact indicates that the social phenomenon has a broader impact on the university community.

H6: Diverse user interests? Figure 14 shows the 20 highest rated courses. We observe that users rate courses ranging from Psychology to Mathematics and Economics. These

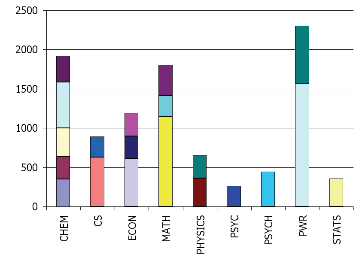

Figure 14: User interests

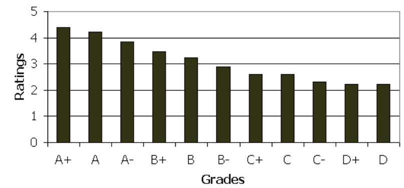

Figure 15: User subjectivity

courses may be required for students following a particular academic program, such as when aiming at a computer science degree. Interestingly, we also observe courses that are not required for any particular degree but are popular choices among students, such as Writing and Rhetoric courses (with code `PWR`).

User Behavior

H7: The objectivity principle? In a community such as a university, one may think that users express informed opinions, and therefore it is easier to trust what users say. Students who post ratings may be regarded as experts who have had significant experience with the courses. They may also have consulted with a number of other students who share the same perspective, so that online ratings may represent a far larger and more representative sample of students than the numbers suggest. Even if some students give highly biased ratings, these may be balanced between those that are positive and those that are negative. A recent analysis of ratings in `ratemyprofessor.com`, a site where college students can evaluate professors, suggests that online ratings in their current form may be useful but possible abuses could limit validity (Otto, Jr, and Ross 2008).

In our system, user ratings range from 1 (lowest rating) to 5 (highest rating). We grouped user ratings based on the grades the users took for the courses they rated: all ratings given for courses where users took an A, all ratings for courses where users took a B, and so forth and we computed the average rating for each grade. Figure 15 shows a strong correlation of average ratings and grades. In fact, user ratings follow user grades. Students who did not perform well in courses gave low ratings to these courses. Consequently, ratings and reviews in a social site, as in real life, may reflect an informed opinion on reality but also often express biased opinions and personal feelings.

H8: Spam is everywhere? General-purpose social sites are open to anyone and users can upload to the system virtually

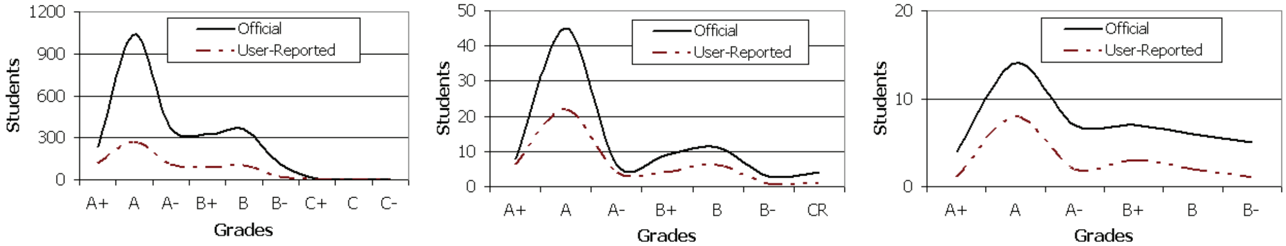

(a) Overall distribution (b) A popular course (c) A less popular course

Figure 16: User honesty: official vs. user-reported grades

anything once they sign up. Users use fake identifiers and there is practically no easy way to verify their true identity. As a result, social sites become susceptible to spam. Spam can take many forms, e.g., as inappropriate comments, malicious posts, misused tags, and so forth.

In CourseRank, users can add comments for courses and they can also flag a comment as inappropriate. The flagged comments are reviewed by an editor, who can delete the comment or leave it in the system. There are 1760 comments in the system. Only 97 comments have been reported as inappropriate, i.e., around 5%, and finally, 20 have been deleted from the editors, i.e., only 1.11% of the comments in the system. CourseRank targets a small, closed community, where users are tied to unique identifiers. We believe that this usage differs greatly from what is seen in general-purpose social sites that are open to anyone and, hence, may attract spammers and malicious users. In CourseRank, users are more willing to contribute more thoughtfully.

The fact that only a small subset of the comments flagged inappropriate have been actually removed from the system highlights another aspect of the problem: subjectivity. What one person may consider as inappropriate content, another person may think it is very appropriate. There are of course behaviors that most people would agree are inappropriate, but defining such behaviors precisely is not easy (Koutrika et al. 2008).

H9: Everyone lies? In many social sites, users often lie when providing information about themselves hiding conveniently behind multiple, fake personas. Many sites use mechanisms to incentivize their users. For instance, (Y! Answers) uses a scoring scheme: providing a best answer is rewarded by 10 points, voting on an answer that becomes the best answer increases the voters score by 1 point, and so forth. Such incentives do not necessarily make users contribute sensibly and honestly.

In CourseRank, students enter their grades for planning their courses. Figure 16(a) shows the correlation of official vs. unofficial grades for engineering courses (we have the official distribution only for these courses.) The self-reported grade distribution follows the same trend with the official grade distribution showing that students are entering valid data. The two distributions do not exactly match because not all students have reported their grades for the courses in the system. Figures 16(b) and 16(c) show the correlation of official and user-reported grades for a popular (i.e., with many students) and a less popular course. We observe that there is a strong correlation between official and unofficial grades both for popular and more obscure courses.

We believe that providing meaningful incentives is very important to make users contribute sensibly. For example, in CourseRank, students provide personal information (e.g., their class, major), the courses they have taken and their grades, because they can use tools, such as the course planner and the calendar, to help them structure their courses over multiple years. For instance, the planner has been an extremely useful feature, so users have a reason to visit beyond just looking for courses to take. It is also a sticky feature. Once a student has taken the time to enter his courses and grades, he keeps returning. The planner motivates the student to enter accurate data: since it shows to its owner grade averages per quarter, and missing requirements for graduation, there is little reason to lie about courses taken. Users are honest, because, in this case, honesty pays off.

Conclusions

In this paper, we presented CourseRank, an educational and social site where Stanford students can explore course offerings and plan their academic program. We analyzed several aspects of user interactions and user-contributed content in the site, such as activity levels, user behavior and user content quality, based on CourseRank data including user contributed information and user logs. Our findings provide useful insights and show the potential of closed-community social sites like CourseRank. These sites can offer valuable, higher-quality, services to closed communities (e.g., academic, corporate or scientific communities) that go beyond simple resource sharing as in general-purpose social sites on the Web.

We see the following important ingredients to the success of special-purpose, closed-community social sites.

1. Added-value services.

The value of CourseRank lies in helping students organize their academic program. Hence, it is not used just as a hobby but it helps them with their work. The provision of a valuable service makes users more actively engaged and more thoughtful contributors.

2. High-quality data.

A social site needs interesting high-quality data. While some of the data in CourseRank is entered by users (course evaluations, courses taken, self reported grades), one key to its success was the availability of useful external data, such as course descriptions and schedules, associated textbooks, official grade distributions, and so on. Having official data in combination with user input adds value to the system.

3. The community feeling.

An important feature of CourseRank is its closed community. Closed communities tend to breed more active users because the feeling of the community membership is stronger. Furthermore, the existing community structure - unlike online communities on the open web - helps influence people to become more willing contributors. For example, students encourage other students to use CourseRank. In a corporate environment, if several members of a group share information through the corporate social site, then the other members will eventually follow.

4. Meaningful incentives.

In a social site, there need to be incentives for users to visit and to share their resources. The incentives are especially critical in the early stages, where there are few resources shared by others and few users. Our Question and Answer Forum provides an example of low traffic caused by lack of incentives to visit and contribute.

Providing meaningful incentives is very important for another reason too. Users may contribute more sensibly and honestly. Our course planner provides a good example. The planner motivates the student to enter accurate data: since it shows to its owner grade averages per quarter, and missing requirements for graduation, there is little reason to lie about courses taken.

References

Adamic, L.; Zhang, J.; Bakshy, E.; and Ackerman, M. S. 2008. Knowledge sharing and yahoo answers: Everyone knows something. In *Proc. of the Int'l WWW Conf.*

Adar, E.; Weld, D.; Bershad, B.; and Gribble, S. 2007. Why we search: Visualizing and predicting user behavior. In *Proc. of the Int'l WWW Conf.*

CiteULike. http://www.citeulike.org/.

Coelho, P. 2008. Speech for the opening of the frankfurt book fair. *http://paulocoelhoblog.com/wp-content/uploads/2008/10/FrankfurtBookFair.pdf*.

Del.icio.us. http://del.icio.us/.

Flickr. http://www.flickr.com/.

Golder, S., and Huberman, B. A. 2006. Usage patterns of collaborative tagging systems. *Journal of Information Science* 32(2):198–208.

Gyöngyi, Z.; Koutrika, G.; Pedersen, J.; and Garcia-Molina, H. 2008. Questioning Yahoo! Answers. In 1^{st} *Workshop on Question Answering on the Web*.

Horowitz, B. 2006. Blog: http://www.elatable.com/blog/?p=5.

Koutrika, G.; Effendi, F.; Gyöngyi, Z.; Heymann, P.; and Garcia-Molina, H. 2008. Combating spam in tagging systems: An evaluation. *ACM TWEB* 2(4).

Lampe, C.; Ellison, N.; and Steinfield, C. 2008. Changes in use and perception of Facebook. In *Proc. of the Int'l CSCW Conf.*

Marlow, C.; Naaman, M.; Boyd, D.; and Davis, M. 2006. Position paper, tagging, taxonomy, flickr, article, toread. In *Hypertext*, 31–40.

Millen, D.; Feinberg, J.; and Kerr, B. 2006. Dogear: Social bookmarking in the enterprise. In *Proc. of CHI Conf.*

Nielsen, J. 2006. Alertbox: http://www.useit.com /alertbox/participation_inequality.html.

Otto, J.; Jr, D. S.; and Ross, D. 2008. Does ratemyprofessor.com really rate my professor? *Assessment and Evaluation in Higher Education* 33(4):355 – 368.

Sen, S.; Lam, S.; Rashid, A.; Cosley, D.; Frankowski, D.; Osterhouse, J.; Harper, F. M.; and Riedl, J. 2006. Tagging, communities, vocabulary, evolution. In *CSCW'06*.

White, R., and Drucker, S. 2007. Investigating behavioral variability in web search. In *Proc. of the Int'l WWW Conf.*

Xu, Z.; Fu, Y.; Mao, J.; and Su, D. 2006. Towards the semantic web: Collaborative tag suggestions. In *WWW Collaborative Web Tagging Workshop*.

Y! Answers. Scoring system. *http://answers.yahoo.com/info/scoring_system*.