

### 3rd Workshop on Online Social Networks (WOSN 2010)

June 22, 2010  
Boston, MA

#### **SESSION 1**

*Summarized by Saptarshi Ghosh (saptarshi.ghosh@gmail.com)*

Balachander Krishnamurthy opened WOSN 2010 by welcoming the attendees and thanking the contributors. He read out welcome messages from Program Chairs Bruce Maggs and Andrew Tomkins (who could not be present).

- ***Ghostbusting Facebook: Detecting and Characterizing Phantom Profiles in Online Social Gaming Applications***

*Atif Nazir, Saqib Raza, Chen-Nee Chuah, and Burkhard Schipper, University of California, Davis*

Atif Nazir presented techniques to characterize and detect phantom profiles in online social gaming applications. He highlighted the fact that social gaming, platforms for which are provided by several popular OSNs, is now a billion-dollar industry, yet no one has previously studied the impact of social games on the underlying social graph. Highly engaging social games provide a tendency for some gamers to cheat by creating fake (phantom) profiles which contaminate the social graph. Whereas most OSNs presently rely on reports by users and manual inspection to detect phantom profiles, the objective of this work is to characterize phantom profiles in a gaming application and devise a supervised algorithm to detect such profiles.

The authors study the phantom profiles created in a social game played in Facebook, the “Fighters’ Club” (FC) game, where two users start a fight and friends of either user can support them. A set of 13 OSN-and-game-biased attributes were tested to identify phantom profiles. However, the social-network-based properties (e.g., number of friends,

number of networks joined, cumulative distribution of number of friends) were found to be similar for the genuine profiles as well as phantom profiles; hence statistical classification using such attributes are not effective in detecting phantom profiles, and game-based properties must be used.

Nazir reported observations on some of the game-based attributes: total number of fights defended is more for genuine profiles than for phantom profiles; average number of opponents in picked fights is observed to be higher for phantom users than for genuine users; phantom users instigate/defend fights of smaller duration. However, none of these attributes individually differentiates well between phantom and genuine profiles; so the authors propose to combine these attributes using a support vector machine classifier. Though the proposed technique achieves impressive results, the speaker's opinion is that the method is not yet perfect; it can be used for, at most, flagging of suspected phantom profiles, but not for automatic deletion of those profiles. In particular, one drawback this method suffers from is that the actual ratio of phantom profiles to genuine profiles in the entire population is unknown.

Someone asked whether a new "arms-race" between creators and detectors of phantom profiles could begin in the near future. The speaker's opinion was that only 5% of the phantom profiles are usually active, and out of these 5%, most will not likely make the effort to create phantom profiles if they find that effective methods for detecting phantom profiles are being used by the OSNs. A member of the audience drew attention to a similar paper on the use of structural analysis to detect malicious user profiles in eBay. Another member of the audience asked whether phantom profiles make social games interesting in some way, but Nazir did not think so.

- ***Diffusion Dynamics of Games on Online Social Networks***  
*Xiao Wei, Jiang Yang, and Lada A. Adamic, University of Michigan, Ann Arbor; Ricardo Matsumura de Araujo, Federal University of Pelotas, Brazil; Manu Rekhi, LOLapps*

In this presentation, Lada A. Adamic discussed propagation of games through invitations sent by users of OSNs to others. In contrast to the previous presentation, this work assumed all profiles to be genuine. The emphasis of this work is to analyze how social games can be designed to propagate efficiently along an OSN, i.e., how more people can be made to play a game. Also, the behavior patterns of gamers in sending out invitations to other users in the OSN were studied. A user is said to be a successful inviter if the user whom he invites actually joins the game. Some of the questions investigated in this work are (1) is there any correlation between how successful one is in inviting friends (to play the game) and how successful the friends are in inviting others? (2) should one invite all one's friends or only a few close friends? (3) can successful inviters be identified based on their profile attributes?

The data of two popular social games were studied using data provided by LOLapps (the company that created the games), both games having millions of Facebook users. They found that most users invite only a few friends; those who invite indiscriminately have a lower yield per invitation. Correspondingly, pacing the invites, sending repeat invitations to the same target, and inviting fewer friends at a time corresponds to successful invitations. On the other hand, factors such as the gender of the inviter or the size of the inviter's friend-network were uncorrelated; however, higher success in invitations is positively correlated with higher engagement (i.e., time spent) in the game of the inviter. Furthermore, dense cliques forming the underlying social network are readily absorbed into the games.

The properties of invitation cascades generated by the invitations sent by users were found to be wide and shallow and to have small-world properties, with many users receiving multiple invitations. It was also observed that the probability of a user joining a game depends on how many invitations they receive but not on how different the inviters are from one another. Adamic pointed out some future directions of research, such as investigating whether the games merely use the underlying social network to propagate or whether the games also cause the social network to grow. Other directions included characterizing large-scale invitation cascades.

Do users who send out more invitations get a higher yield in absolute numbers, even if the per-invite success is lower? Adamic said that this was true, but that games must take into account how much they annoy potential users, which could adversely affect adoption of future games. A member of the audience was of the opinion that most users of OSNs are likely to not use their official profiles for playing such games and sending invites; rather, they would create profiles specifically for the purpose of gaming. Another attendee asked whether the cliques that get absorbed into the game were cliques in the real world. Adamic said this was likely, to which the questioner asked about the profile similarity within cliques. Adamic replied that this would be an interesting measurement, since their work so far had looked at the correlation of profile similarity of inviters with invitation success, not within-clique profile similarity.

- ***Outtweeting the Twitterers—Predicting Information Cascades in Microblogs***  
*Wojciech Galuba and Karl Aberer, EPFL; Dipanjan Chakraborty, IBM Research, India; Zoran Despotovic and Wolfgang Kellerer, DOCOMO Euro-Labs*

Wojciech Galuba presented this study on the characteristics of information flows in Twitter. The authors report that the distribution of number of tweets received by Twitter users has a median of 552 tweets per day, which indicates an information overload. The motivation of the present work is to improve how information flows in Twitter. The authors look into the diffusion of URLs in Twitter (through tweets)

and investigate whether it can predict with what probability a given URL will be mentioned (tweeted) by a given user. Such predictions can provide protection from information overload, e.g., by sorting incoming URLs by probability of re-tweeting.

Data was collected by querying the Twitter search API for the string “http”. Power law distributions were observed for user activity, whereas log-normal fits were observed for the degree distribution of Twitter users in the social network. The authors studied information cascade digraphs formed in Twitter, where nodes in the digraphs represent users who mention a given URL and directed edges indicate the flow of a URL from one user to another through a tweet. Two types of information cascades were observed: re-tweet cascades, which are trees in their structure, and follower cascades, which are directed acyclic graphs. Information cascades were also found to be shallow in their depth.

The authors also proposed a model to predict the probability of a given user mentioning a given URL. The model considers three factors for the prediction: influence of one user over another, external influences on users, and the vitality of the given URL. The input to the model is a time window of tweets, and the outputs are the probability values as stated above; the model optimizes the F-score (the harmonic mean of precision and recall). The proposed model was able to predict more than half of the URL-tweets with only 15% false positives. It is to be noted that the events predicted are only within one hop in the Twitter follower graph from the users that have already mentioned a URL.

Was spam considered in the model? Galuba said that they relied on Twitter itself to remove spammers, and each user in the data was verified later to ascertain whether the user profile still existed. Galuba also said that about one-fifteenth of the tweets gathered in the experiments had tweets in them. Was the presence of private profiles considered in the model? Galuba clarified that it was not considered. What is the relative importance of the three factors considered in the model (stated above)? Galuba’s opinion was that the user-to-user influence seems the most important.

## SESSION 2

*Summarized by Saptarshi Ghosh (saptarshi.ghosh@gmail.com)*

### ■ **Privacy Leakage in Mobile Online Social Networks**

*Balachander Krishnamurthy, AT&T Labs—Research; Craig E. Wills, Worcester Polytechnic Institute*

This work, presented by Craig E. Wills, is a continuation of the authors’ work presented in the 2nd WOSN in 2009. The objective was to study the potential leakage of personally identifiable information from OSNs to third-party aggregators. The widespread use of mobile devices has resulted in the popularity of mobile-OSNs (mOSN); personal information shared by users with a mOSN that is connected to a traditional OSN is also shared with that OSN, and this may

cause additional concerns in leakage of such information. Most mOSNs have a “check-in” mechanism that establishes both a user’s presence in the mOSN and the user’s current location. Also, mobile devices have a unique device identifier and if this identifier is leaked to a third-party aggregator, it can be linked to a user’s real identity. Such threats are usually not present in the case of a traditional OSN being accessed using a browser running on a computer.

The present work examined some mOSNs having roots as a traditional OSN and some OSNs which were not in existence before widespread use of mOSNs. Most of these mOSNs studied have device-specific applications that make the use of the data easier for users (e.g., Apple iPhone applications). However, only a minority of mOSNs provide any privacy settings via the smartphone and mobile applications. Many of the mOSNs studied are connected to the standard OSNs like Twitter, Facebook, and Flickr—comments posted on the mOSNs get reflected on these standard OSNs as well. The OSN identifier of the user posting the comment is often leaked from these mOSNs.

Wills reported that some type of private information, such as user location or device identifier, is being leaked by all the mOSNs that were studied. Such leakages may be quite difficult to prevent by the users. The authors classify information leakages to third parties from mOSNs into two classes—explicit leakage, via Request-URI or POST requests, and implicit leakage, via the Referer or Cookie HTTP headers. Explicit leakages are difficult to prevent unless done on a per-server basis, while implicit leakages can be prevented by users. Moreover, users nowadays want to share their profile information across several OSNs, and this makes some of the personal information available to third-party aggregators.

What causes such information leakage? Is it careless programming, or is it part of the business model? The co-author of the paper, Balachander Krishnamurthy, said that the authors had contacted some of the mOSNs studied to inform them of the leakage and to seek explanations, but no satisfactory explanation was given; rather, one mOSN replied that such leakages are common. Have the authors tried to correlate what is public information in an OSN (i.e., what the users want to make public) to what was being leaked? It has been verified that some OSNs are actually handing over non-public information to third-party aggregators, allowing the real identity of users to be established or the tracking of Web sites visited by users.

### ■ **Don’t Tread on Me: Moderating Access to OSN Data with SpikeStrip**

*Christo Wilson and Alessandra Sala, University of California, Santa Barbara; Joseph Bonneau, Computer Laboratory, University of Cambridge; Robert Zablit and Ben Y. Zhao, University of California, Santa Barbara*

Christo Wilson presented methods to restrict or moderate access to user data in OSNs. User profiles in OSNs include

information (such as birthday, location) that can be used for malicious purposes (e.g., by spammers to construct phishing mails). Researchers also collect and analyze large amounts of OSN data. In addition, there are firms which provide features for the very large-scale crawling of OSNs. OSN users are sometimes unable to defend themselves from such information gathering due to complex privacy settings used by some OSNs; most users are also too lazy to use these settings.

Wilson discussed some of the existing technologies to control crawling, but none of these technologies sufficiently controls access to OSN data by today's crawlers. Configuration files (e.g., robots.txt) tell crawlers how to behave, but compliance with the rules stated is voluntary. Requests can be filtered by servers based on HTTP Request Headers, but headers can be modified by crawlers. Servers also use IP-address-based tracking, but NATs and proxies create problems; for example, many genuine users using the same proxy as a malicious crawler can get blocked. Several OSNs do authenticate user accounts and track the number of requests sent by accounts, but the problem in this approach is that the URLs are session-independent; hence a malicious crawler can switch to another user account if one is blocked due to generation of too many requests.

To effectively moderate the crawling of OSN data, the authors propose SpikeStrip—a technique using server-side link encryption by the user's session key and a secret key generated by the server. If a crawler switches to another account, the session key changes and the server will no longer be able to decrypt URLs sent by the crawler using the new session key (since the encryption was done using the previous session key). Thus this method prevents session-switching by crawlers. Also, distributed crawlers will be prevented from sharing URLs, since the session keys will be different for each instance of the crawler. This method also enables strong rate limiting. The authors evaluate their proposed approach using two metrics—the server overhead caused by the method, and the effectiveness of the method in throttling crawlers. The proposed technique was implemented on an Apache server, and resulted in about 7% additional overhead on the server while very successfully throttling crawlers.

A member of the audience pointed out that a link-encryption-based technique has been commercially used by an OSN in Russia since 2007 to moderate access to data. Whereas the authors of the present work assume that a session key is linked to a user account, the aforementioned OSN server sends back new session keys with every request. Someone also commented that many OSN users wish to share links but link encryption makes this difficult. Wilson's opinion was that crawlers and users (people) are usually interested in accessing different sets of URLs: people are very often interested in accessing photos (images), for example, unlike crawlers; SpikeStrip could be

used to selectively encrypt the links that crawlers are mostly interested in.

- **Prediction Promotes Privacy in Dynamic Social Networks**  
*Smriti Bhagat, Rutgers University; Graham Cormode, Balachander Krishnamurthy, and Divesh Srivastava, AT&T Labs—Research*

This work, presented by Smriti Bhagat, discusses ways to anonymize OSN data in order to protect the identity and privacy of the users before the data is published for use by researchers and other third parties. For anonymization, a version of the OSN data must be published which protects the identity of users, yet is usable for analysis (e.g., studying communication patterns). Some commonly used techniques for anonymization are simple anonymization (replacing user names with arbitrary numbers, but easy to break), adding/removing edges to make neighborhoods similar (but this alters the social graph), and node mapping (grouping multiple users to hide the identity of any individual user). However, the focus till now has mostly been on anonymization of a static network, publishing a single anonymized instance of the network. The present work is on the problem of dynamic network anonymization—given a temporal sequence of graphs  $\{G_0, G_1, \dots, G_t\}$ , the output needs to be a (modified) sequence of graphs such that the likelihood of identifying the node that represents a given user or determining if a pair of users interact is small.

Bhagat clarified that naive approaches for dynamic anonymization—such as individually anonymizing each graph instance in the given sequence, or anonymizing only the first instance ( $G_0$ ) and inserting new edges as they appear—do not work well. The proposed method uses node mapping along with link prediction techniques; two users are grouped together only if there is no path of length two with at most one predicted edge between these two users. In simple terms, if there is a high probability of two users getting linked with each other in future, the proposed approach will keep these users in different groups. The authors defined a new metric called Edge Identification (EI) to measure privacy loss due to anonymization and show significant improvement in EI on using link prediction for social network datasets. The authors also showed that graph sequences anonymized by this approach effectively answer queries with insignificant increase in errors due to grouping with link prediction.

Is the proposed technique dependent on the window length? Bhagat clarified that the window length needs to be such that consecutive samples in the given sequence differ considerably with respect to number of new users and edges introduced in the samples. If certain structures get introduced in the social graph with time, would it be possible to identify the structures in the anonymized sequence? Although such structures might be identified in the anonymized sequence, due to the guarantees of the grouping approach, even if a few users in a group are identified, the privacy of other users in the group may be preserved. How-

ever, it is not possible to provide absolute guarantees on anonymization without assuming some attack models and background information.

### SESSION 3

Summarized by Christo Wilson ([bowlinearl@gmail.com](mailto:bowlinearl@gmail.com))

#### ■ **A Geometric Model for On-line Social Networks**

Anthony Bonato, Ryerson University; Jeannette Janssen, Dalhousie University; Pawel Pralat, West Virginia University

Anthony Bonato presented work that focuses on creating new graph models specially tailored to social networks. Typical random graph models are targeted to observed properties of random graphs, such as power-law degree scaling, small-world clustering, and local communities. It has also been observed that dynamic random graphs tend to exhibit densification power law (i.e., average degree increases over time) and shrinking distances (i.e., graph diameter decreases over time). The goal of this work is to create a good model for social graphs specifically, since few models are tailored specifically for this purpose. This model should evolve in a “natural” way, so that generative forces in the social network can be understood and used to predict links in the graph. The authors’ previous work focused on a deterministic model exploiting transitivity in social relationships, called iterated local transitivity. This model exhibits many desirable properties but fails to generate power-law degree distribution. Hence, the authors moved on to geometric models that map OSNs, which live in social space, to Euclidean space. This process involves isolating user attributes that can be used to define a set of dimensions and then determining user placement in the space by looking at the commonalities between these elements. Once users have been embedded in Euclidean space, a volume can be drawn around them, and all other users inside this volume are linked to the central user.

The authors’ current work on this topic involves combining two existing models, random geometric and Protean, to construct a new model called Geometric Protean (GEO-P). The key features of this model involve varying the size of a user’s influence sphere based on a ranking function, iteratively removing and adding new nodes, then re-ranking. The benefits of this model are that it generates power-law graphs where average degree, diameter, and density are fully parametrized. The authors also note that the dimensionality of social graphs appears to be equal to  $\log(n)$ , although this has not been proven yet.

Questions from the audience mostly focused on the dimensionality of social graphs. Specifically, is there a way to map between the Euclidean dimensions and specific user attributes? As the authors’ work is more theoretical in nature, they were unable to concretely answer these questions, instead saying that it’s up to empiricists to solve those types of questions.

#### ■ **Distance Matters: Geo-social Metrics for Online Social Networks**

Salvatore Scellato and Cecilia Mascolo, Computer Laboratory, University of Cambridge; Mirco Musolesi, University of St. Andrews; Vito Latora, Università di Catania and INFN

Salvatore Scellato presented a first stab at answering the question, what is the role of geography in the structure of OSNs, and how does it affect information dissemination? To answer this question, the authors gathered data from several social networks and mapped people based on their stated location. They could then examine users’ friends in geographic space, to examine whether friendships occur between users who are close geographically, and classify users based on their preference for short- or long-distance relationships. The authors also aim to examine information dissemination with respect to user geography and how this may impact social application performance.

The authors constructed a new social graph type called a Geographic Social Network, where all users are placed in a 2D space, and social edges between users are weighted based on geographic distance. This leads to the invention of two new graph metrics: node locality (i.e., how close in space my neighbors are) and geographic clustering coefficient. These metrics are constructed to highlight nodes and groups of nodes that exhibit tight geographic clustering, and they include a scaling parameter to account for relative differences in geography (i.e., comparing user distances in a tight urban metropolis is not equivalent to the same task in a widespread, rural community).

The authors crawled Brightkite completely, and crawled snowball samples of FourSquare, LiveJournal, and Twitter. They then used the Yahoo Geocoding API to obtain GPS coordinates for users who gave their location via textual descriptions. The authors observe that users in location-based OSNs such as FourSquare and Brightkite have friends links geographically closer (1,200 km on average) than more general OSNs (Twitter was 5,000 km on average). On FourSquare, in particular, 60% of links are less than 1 km in length, showing that users tend to live very close to their OSN friends. There even exist a small number of hyperlocal users on FourSquare and Brightkite that have geographic clustering coefficients of 1, meaning they all live in exactly the same place.

Questions from the audience fell into two areas. The first concerned bias in the data due to people not uniformly tagging data (on Twitter) and people having geographic biases (i.e., users tend to follow people in their own country/culture). The authors agreed that these were difficult problems to deal with. The second area concerned time-based dynamics: does geography drive new friendships or do existing friends tend to migrate close together? Unfortunately, no location-based OSN reports the times that new friendship links form, and thus this can’t be examined as of yet.

- **Orion: Shortest Path Estimation for Large Social Graphs**  
Xiaohan Zhao, Alessandra Sala, Christo Wilson, Haitao Zheng, and Ben Y. Zhao, University of California, Santa Barbara

Xiaohan Zhao presented this work, focused on the creation of a new system that can perform node distance computations on massive social graphs in real time. There are many interesting and novel social applications that are enabled by today's huge social graphs. Examples include influence maximization for marketing campaigns, and social search. However, both of these examples, and many others, are predicated on being able to compute social distances in real time. This is impossible with today's graph distance algorithms, which usually run in  $O(VE)$  time.

The goal of this work is to design a new system that can scalably provide social distance information in  $O(1)$  time to enable real-time applications. Meeting this goal requires that some caveats be managed: specifically, the preprocessing time necessary to bootstrap the system, and the accuracy of generated results. The authors proposed to meet these goals through the construction of a novel Graph Coordinate System called Orion that embeds users from the social graph into a Euclidean space.

There are two approaches to embedding social graphs into Euclidean space. The first approach, modeling a physical system, requires too much preprocessing to be practical. The second approach, using landmarks, works better in practice. High-degree users are the best choice for the landmarks, and all other users can be placed by performing a constrained number of full-breadth first searches of the graph. The authors use an incremental global simplex downhill heuristic to ensure that preprocessing time for Orion remains linear on the size of the graph.

The authors evaluated Orion using four datasets crawled from Facebook regional networks. They show that a 300k node graph can be preprocessed in about two hours (an acceptable, one-time cost). They also showed that Orion calculates node distance measurements seven orders of magnitude faster than BFS (Breadth-first search), with only 0.2 relative error in distance calculations. The authors also show that Orion is suitable for calculating higher-order graph metrics, such as average path length (error  $> 0.3$ ).

Questions from the audience were varied. One person wanted to know about recovering actual, hop-by-hop shortest paths from Orion rather than just overall distances, which the system doesn't currently support. Another question concerned dynamic graphs, which the presenter agreed was an interesting problem and the likely future focus of this work.

- **The Effects of Restrictions on Number of Connections in OSNs: A Case-Study on Twitter**  
Saptarshi Ghosh, Gautam Korlam, and Niloy Ganguly, IIT Kharagpur, India

Saptarshi Ghosh discussed the effects of the limits imposed by many OSNs on the number of social friends each user may have. These limits are imposed mainly to control spam

and to prevent undue strain on the OSN's infrastructure in supporting user-to-all-friends communication. In contrast to fixed limits in most OSNs, Twitter has imposed a "soft" limit on the number of people a user can follow (i.e., users' outgoing links are limited based on how many incoming links they have). While people often complain about these limits, research results suggest that the limits only actually affect hyper-active users. However, these limits do have ramifications on how the OSN evolves. This work seeks to probe those ramifications.

On Twitter, there is the concept of "follow spam," i.e., users who do nothing but follow others just to draw attention to themselves for nefarious reasons. This causes Twitter to limit the number of people one can follow to 2000, beyond which one can follow more only if one has a commensurate number of followers. The actual algorithm for this is secret (security through obscurity), but people have independently reverse-engineered two plausible algorithms. The authors performed BFS crawls of 1 million Twitter users in November 2009 (about one year after the limit was imposed) to observe users' in-degree and out-degree distributions. They observed that only 6.6% of users follow more than 2000 people, and almost all users have out-degree less than 1.1 times their in-degree; this result agrees with the conjectures regarding Twitter's secret follow-control algorithm. The authors also note that the 2000 follow-limit on out-degree causes a pronounced spike around 2000 in the out-degree distribution, as there are many users who exhibit the "follow spam" paradigm of following many but not being followed in return. The authors discuss a new social network growth model based on preferential attachment, which incorporates restrictions on user-degree. The proposed model can be used to analyze effects of different forms of restrictions on the topological properties of social networks as well as to design restrictions of varying rigidity.

Questions from the audience focused on potential deficiencies in the authors' crawling methodology. Specifically, BFS (and repeated BFSes from distinct starting points) are likely to bias gathered data towards high-degree nodes. The speaker agreed with this observation and said that they plan to collect unbiased samples of the Twitter social network, using methods such as Metropolis-Hasting random walk.

## SESSION 4

Summarized by Christo Wilson ([bowlinearl@gmail.com](mailto:bowlinearl@gmail.com))

- **Voice of the Customers: Mining Online Customer Reviews for Product Feature-based Ranking**  
Kunpeng Zhang, Ramanathan Narayanan, and Alok Choudhary, Northwestern University

This work, presented by Kunpeng Zhang, focuses on aggregating and automatically distilling textual user reviews of products online into concise and comparable feature-based rankings. Online shopping continues to increase its market share, and many of the top sites encourage users to review

products they have purchased. User reviews are very important in influencing the decisions of new buyers, but it can be time-consuming to read all available reviews, especially since different people have different requirements when evaluating products. In this work, the authors propose using comparative and subjective evaluative sentences to construct graphs of product features which are weighted according to the products' relative strengths and weaknesses.

The authors look for two different types of statements when parsing user reviews: comparative sentences and subjective sentences. The first type of statement compares the relative merits of one product as compared to one or more other products in the same category. These statements form the edges of the product graph, with the weight being the positive or negative customer sentiment. The second type of statements are simply positive or negative declarations about one specific product. These statements are used to weight nodes of the product graph. The authors leverage a corpus of 2000 positive keywords, 4500 negative words, and 30 words of negation when performing their textual analysis. The authors use three techniques (keyword matching, part-of-speech analysis, and predefined patterns) to identify comparative sentences, at which point they apply refinements to ensure that statements conform to the requirements necessary for further analysis.

The authors targeted this work towards evaluation of cameras and televisions using data crawled from Amazon. Each product is present in multiple different weighted graphs, where each graph represents a particular product feature, i.e., a camera might be in lens, flash, and battery-life graphs. Once these graphs are constructed the authors leverage a novel ranking algorithm, pRank, to calculate the importance of each product  $n$  for each feature  $f$ . This allows them to isolate the most important feature in terms of the overall product quality (importance function). The statistics of feature sentences also tell what features were most talked about (relative feature fraction). pRank compares favorably with reviews generated by experts in each of the two measured product categories, as opposed to simple metrics like Customer Star Rating and SalesRank, which fail to correlate.

Questions from the audience focused on extensions to this work. Specifically, could the weight of individual reviews be added to the system, say to up-moderate knowledgeable experts and down-moderate trolls? The author stated that they are currently working on using data on reviewer fidelity from Amazon to extend their techniques to incorporate this additional information.

■ ***I rate you. You rate me. Should we do so publicly?***

*Chun-Yuen Teng, Debra Lauterbach, and Lada A. Adamic, University of Michigan*

Chun-Yuen Teng presented this paper, which focused on quantifying the biasing effects of public visibility of user reviews in online systems. Many Web 2.0 sites use recommendation/reputation systems to help users evaluate products and each other, but are these social ratings honest? Specifically, does the design of these sites, and whether

review sources are public or private, affect the quality of resultant reviews? To answer these questions, the authors crawled reviewer data from Amazon and Epinions and obtained the entire (anonymized) review database from CouchSurfing.

The most immediate effects of site design observed by the authors were on review reciprocity. Positive reviews left in public places were likely to generate positive reviews for the initial reviewer as well. However, making reviews/reviewers public (and thus leaving open the threat of public retaliation) also has the effect of causing few overall negative reviews to be left at all. For example, the ratio of positive to negative reviews on CouchSurfing is 2500:1, and 70% of vouches are reciprocated. The same is true on Epinions, where public reviews are likely to be positive, while anonymous reviews are more negative. Real name (as opposed to anonymous) reviewers also tend to be more thorough, as these reviewers leave longer comments on Amazon.

Other social factors seem to affect reviews as well. Men tend to be more egalitarian with their ratings, while women tend to rate other women more highly on friendship and trust. Age, geographical separation between reviewers, and cultural homophily also all influence review scores between people. Overall, the authors' takeaway advice was not to take online reviews at face value; there are many hidden social effects that warp and skew reviews.

Questions from the audience focused on additional sources of bias in the review data. The authors were asked if temporal dynamics affect reviews (i.e., incidental mood swings) and about review reproducibility (i.e., if one person reviews the same thing twice, will the reviews be consistent?). The authors acknowledged that these were difficult questions that were hard to quantify with the available data.

■ ***Measuring Online Service Availability Using Twitter***

*Marti Motoyama, University of California, San Diego; Brendan Meeder, Carnegie Mellon University; Kirill Levchenko, Geoffrey M. Voelker, and Stefan Savage, University of California, San Diego*

Marti Motoyama explained the potential for using Twitter to automatically detect outages affecting online service providers. Twitter is ideal for this sort of study because of the social structure of the network, the real-time nature of its messages, and the rapid dissemination of information throughout the community. People also provide a varying set of real-world vantage points on service conditions and may notice localized and transient failures that other automated up-time measurement tools (e.g., ping) may miss. In effect, this enables researchers to treat the OSN as a human-sensor network. The importance of measuring service outages will only increase as more services move to the cloud; hence the need for this work. As a motivating example, the authors note that it took several traditional news media outlets over one hour to publicize news of the Gmail outage that occurred on September 1, 2009, whereas users on Twitter started messaging about the outage within minutes of its occurrence.

The authors crawled Twitter for data and performed textual analysis to look for outage events. They focused on the bi-gram “is down” as the primary indicator of service outages, with the hashtag “#fail” being a secondary indicator (mostly because it is a commonly used tag, hence leading to a noisy signal). They examine the two words surrounding service names during known outages and verify that the word preceding the phrase “is down” effectively denotes the name of the affected service. In order to filter out noise (e.g., people tweeting about non-existent outages), the authors use an exponentially weighted moving average to determine the incidental volume of tweets about each service. The stream of tweets is divided into five-minute intervals, and any service that exceeds a moving threshold for two consecutive intervals generates an outage alarm. In order to tune the parameters, the authors created a validation set by looking for outages through search engines and manually inspecting maintenance blogs. Subsequently, they tried different combinations of parameters and picked the ones that produced the lowest F-scores.

In order to validate their methodology, the authors looked at eight service disruptions that occurred in 2009, for which they were also able to find an article or blogpost approximating the start time of the outage. Their automated system managed to detect all eight test events, although the time to detect the event sometimes lagged 10 to 50 minutes behind the actual start of the outage. They noted that news articles may be imprecise in their assessment of the actual outage start times. The authors observed that detection times could be improved by expanding the set of warning bi-grams, e.g., to detect tweets such as “anyone having problems with Gmail” instead of just “Gmail is down.”

Running the automated analysis on the entire corpus of crawled tweets resulted in the detection of 894 outage events affecting 245 entities. The authors determined that of these 245 entities, 59 were false positives, mostly associated with sporting events (where the phrase “is down” makes sense). Of the top 50 experimentally discovered outages (as determined by the volume of tweets containing “is down”), the authors manually verified 48. Interestingly, nine of these outages were of Twitter itself. One possible explanation may be that third-party applications queue tweets during the outages and push the updates out when Twitter comes back online. Out of 50 random outages sampled by the authors, 35 were manually verifiable. These included outages of major services such as World of Warcraft and Netflix.

One person was concerned about using the system to measure public sentiment, which Motoyama agreed would be possible. Another question was about attacks, i.e., would it be possible to spam/astroturf Twitter and fool the sensor? This seems unlikely, given the large number of users on Twitter, but is not impossible. Lastly, there was a question about using tweets that are geo-tagged to further refine the locations of failures. The author thought this would be a very valuable feature, but presently this isn't feasible given the small percentage of tweets that are geo-tagged.