

Reality Mining and Personal Privacy

Will Privacy Disappear when Social Sensors Learn Our Lives?

Anmol Madan, Benjamin N. Waber, Margaret Ding, Paul Kominers, and Alex (Sandy) Pentland

MIT Media Laboratory

20 Ames St.

Cambridge, MA USA 02142

{anmol, bwaber, mjding, pkoms, pentland}@mit.edu

Abstract— We define reality mining as quantifying and modeling long-term human behavior and social interactions, by using mobile phones and wearable badges as sensors that capture real-world face-to-face interactions. In this paper, we describe two experiments that use this approach: (a) understanding the diffusion of social behaviors using mobile phones at an undergraduate dormitory for an entire year and (b) studying the interactions that influence workplace and task-level productivity at a data server configuration firm, using badges for an entire month. We illustrate that this approach allows us to build better computational and social models of human behavior.

Such data and tools are promising for computational social science applications, but the legal and ethical boundaries around data ownership and user privacy are still unclear. For example, who owns such employee data in a corporate setting? In the consumer setting, what data access rights should service providers have? How are users who don't directly participate in research or use applications affected? In the second part of this paper, we discuss our perspective on these questions within the context of the above experiments.

Keywords- human behavior modeling; reality mining; privacy and data security

I. INTRODUCTION: REALITY MINING

As citizens of the information age, we leave pervasive digital trails that characterize our behaviors everywhere – in emails, on online social networking sites, in mobile phone call logs, in ATM machines, in metropolitan train systems. Researchers from different quantitative fields are leveraging extensive interaction data collected using mobile phones, wearable sensors and various online social tools to build underlying theories about human behavior. In the past, social scientists have relied on survey instruments to capture such interaction data. However, surveys do not provide fine-grained data about the user's day-to-day interactions or communication with others. In addition, human errors are induced into surveys due to time omission (i.e. the memory of events and actions decays with time), telescoping effects (i.e. individuals tend to under-estimate the time dimension of an event) and selective memory bias (e.g. people find it easier to remember social interactions with people they admire or are

attracted towards). In a survey of informant accuracy literature, Bernard and colleagues found that recall of social interactions in surveys is typically in the range of 50% accuracy [4].

Mass-market mobile phones are pervasive, long-term social sensors. Eagle and Pentland [9] coined the term Reality Mining, and used mobile phone Bluetooth transceivers, phone communication logs, and cellular tower identifiers to discover the social network structure, recognize social patterns in daily user activity, infer relationships, identify socially significant locations, and model organizational rhythms. Gonzalez et. al [11] analyzed GPS location traces for 100,000 individuals and found that a simple spatial probability distribution could be used to characterize human mobility patterns better than random walk or Levy flight models. Onnela and colleagues [21] used phone communication logs to characterize the local and global structure of a 4.6 million-node network, and found that intermediate strength ties play a key role in the diffusion of information.

Similarly, socially aware electronic sensor badges have also been used to capture interactions and learn the structure of social networks. Choudhury and Pentland [6] designed the Sociometer, a wearable sensor package for measuring face-to-face interaction between people using an infrared (IR) transceiver, a microphone and accelerometers. Face-to-face interactions captured using the Sociometer were used to model the structure and dynamics of social networks. The Sociometric badge [20] was designed to identify human activity patterns, analyze conversational prosody features and wirelessly communicate with radio base-stations and mobile phones. Sensor data from these badges has been used in various organizational contexts to automatically predict employees' self-assessment of job satisfaction and quality of interactions [26]. Other researchers have used online interaction data from emails, viral recommendations or social networking sites. [1, 18] However, in this paper, we focus on tools that capture face-to-face interactions.

In the following sections, we provide two examples of reality-mining experiments and behavioral inferences. In the

final section, we discuss the privacy implications of our research.

II. FIRST STUDY: MODELING SOCIAL DIFFUSION USING MOBILE PHONES

A. Research Goal and Experiment Design

Social networks play a fundamental role in the propagation of ideas, opinions, innovations, recommendations and media. *Diffusion* is the phenomena of propagation within a social network. *Social influence* is the ability of a node to manipulate the propagation process, by inducing other nodes to adopt or reject the transmission. Models of social diffusion and influence have been studied in many different forms, i.e. the transmission of political opinions and news in political science [15]; the diffusion of innovations in management science [24]; the value of novel information in organizational behavior [1]. Several simple probabilistic models of diffusion processes, like the threshold model [12] have been proposed.

In order to create realistic predictive models of diffusion phenomena, we need to train with a complete picture of the social interactions between participants and the exogenous variables that affect the transmission process. An important aspect missing from prior work is fine-grained data about communication and face-to-face interaction between individuals. With mobile devices that capture face-to-face interaction, we can explore questions like-- if we measure who talks to whom, and how often, does that represent the transmission probability between two people? Does regular co-location or frequent communication imply greater social influence? What is the role of different types of communication and interactions, e.g. the interaction in the workplace or in social milieu – do they translate into different types of social influence? Is one type of interaction more powerful than the other?

To understand these diffusion behaviors, we outfitted sixty-five undergraduate residents of a university dormitory with Windows Mobile smart-phones, enhanced with software for long-term data collection. These participants represented eighty percent of the total population of the dormitory, which is known for its pro-technology orientation and tight-knit community.

- The phones periodically scanned for Bluetooth wireless devices in proximity. Mobile phones are equipped with class 2 Bluetooth radio transceivers, which have a maximum range of 10 meters. Bluetooth and other wireless-radio based co-location techniques have been used to identify the nodes and edges in the social network graph [11].
- The phones periodically scanned for Wi-Fi (WLAN 802.11b) access point identifiers. Since the university campus has high Wi-Fi penetration, these identifiers can be used to infer homogeneity and entropy of location and proximity patterns, e.g. is there a cluster of users who tend to visit similar locations frequently?

- All phone call logs and SMS logs were captured. The temporal and frequency features extracted from communication logs can be used to infer strength of social ties and identify relationships, e.g. how often do certain people call on weekends?
- A custom music player was installed on the phone, which allowed participants to play, share, rate and search through the music library. Participants had access to over 1500 independent music tracks from many different genres. All events were logged on the server-side, and user-ratings were used to control for music quality in the analysis. To send a track to any other participant, participants could simply click on the ‘share’ button on the mobile phone application and select the recipient.
- To eliminate confounding effects, special care was taken to ensure that the music was not featured in mass media or was otherwise familiar to the participants. All the content was sourced under the Creative Commons license or with explicit permission from the ‘indie’ artists.

In addition to the music propagation portion of the study, participants completed monthly surveys that helped us model diffusion along the following behavioral dimensions:

- Sociometric survey for relationships with other residents, where possible options were ‘close friend’, ‘casual acquaintance’, or ‘don’t know’
- Monthly surveys about political opinions, health, flu, exercise, stress, smoking and other social behaviors.

B. Analysis and Results

1) Predicting Relationships:

The following features were extracted for every participant dyad and used in the subsequent analysis of relationships and sharing behavior.

- Communication features:
Total communication, off-peak communication (after 11pm and before 8am), weekend communication (Saturday and Sunday of the week), incoming versus outgoing communication and SMS communication
- Location features:
Jensen Shannon divergence between distributions of the first hundred most-frequently observed WLAN IDs between individuals. Co-location based on WLAN ids has low resolution (100-300 feet indoors) and was not used.
- Where mentioned below, the number of music tracks shared between two participants was also used as a feature

To train a model that predicts the relationships between participants using social interaction data, self-assessments of relationships between dyads (‘friend’, ‘acquaintance’, or ‘don’t know’) from the sociometric survey were used as training labels. The communication and location features are correlated with the user-stated relationship ($r = 0.6$, $p < 0.01$). In addition, if we consider that music spreading through a social network is an ‘active probe’ that reflects the strengths of social ties, and

use the number of music tracks shared as an additional feature, this correlation improves ($r = 0.66$, $p < 0.01$). The communication and location features help discriminate between different types of relationships, i.e., friends vs. acquaintances. The total communication and total number of shares between individuals are positively correlated with both friends and acquaintance types of relationships. The off-peak communication and SMS communication features were positively correlated only with the ‘friend’ relationships, and not with the ‘acquaintance’ relationships.

These features can be used to build a Bayesian network classifier (cost-sensitive, 5-fold cross-validation) that predicts whether two individuals are close friends based on the communication data available for them. The overall accuracy of such a classifier using only ‘passive’ mobile phone features is 87.3% (f-measure = 0.646 for the Friends class). Similar to the above case, we can improve the accuracy of such a classifier by using the number of tracks shared between two people as an additional feature to 90.1% correctly predicted relationships overall (f-measure = 0.727 for the Friends class). Since only 28% of all possible dyads are friends, a cost-sensitive approach is used in model training and classification errors for the ‘Friends’ class were penalized more than the ‘not-Friends’ class by a factor of 3.

2) Predicting the Sharing of Music Between Dyads

The communication and location features extracted from mobile phone logs are correlated with observed sharing of music ($r = 0.65$, $p < 0.01$). The specific features that are important predictors of sharing are: total calls and total off-peak duration, SMS communication and the KL divergence of WLAN IDs. Dyadic sharing behaviour shows a higher correlation with automatically captured mobile phone features than self-reported relationships ($r = 0.42$, $p < 0.01$ for mutually acknowledged Friends). This result indicates that social interactions automatically captured using mobile phone sensors may be better predictors of the transmission probability than user self-assessments.

The media propagation observed in the experiment was further broken down into two distinct types:

- Approximately 70% of the total shares were between ‘mutually acknowledged friends’. For this subset of dyads, the correlation of location and communication features with propagation is even higher. This reflects diffusion within cohesive social ties.
- The remaining 30% of shares were between strangers or weak ties. For this subset of dyads, the location and communication features are not significantly correlated with sharing. This form of diffusion is consistent with the theory of weak ties.

The observations of sharing between participants can be broken into a 2-class (sharing /no-sharing) or 3-class model (‘no sharing’; ‘low sharing’; and ‘high sharing’; class boundaries were selected based on the distribution of shares). Without any prior relationship data and using only mobile phone features, the 2-class prediction accuracy using a cost-sensitive Bayesian network classifier is 71.5 % (precision = 0.69, recall = 0.426 Sharing class). With a similar model, the 3-class, 5-fold cross-validation accuracy is 69%. By implementing a hierarchical Bayesian model, where

relationships are inferred from mobile phone features, the 2-class classification accuracy for sharing increases to 74%.

III. SECOND STUDY: MODELING INTERACTIONS AT THE WORKPLACE USING BADGES

A. Goals and Experiment Design

Studying organizational behavior in detail over long periods of time has long been a challenge to the social science community [1,3,7]. Human observers are expensive, suffer from subjective opinions, and it is difficult for them to remain unobtrusive in an organizational environment. As described previously, surveys based on participant recall suffer from memory effects. More recently, e-mail and other forms of electronic communication have been employed to examine relationship structure (i.e. social network structure) [13]. This research has led to a greater understanding of how organizations function and what management practices lead to greater productivity, but important communications are usually face-to-face [17].

What is necessary to alleviate these problems is a device that could automatically record the behavior of hundreds of individuals with high accuracy over long periods of time [22]. We have created a wearable Sociometric badge that has advanced sensing, processing, and feedback capabilities [20]. In particular, the badge is capable of:

- Recognizing common daily human activities (such as sitting, standing, walking, and running) in real time using a 3-axis accelerometer
- Extracting speech features in real time to capture nonlinguistic social signals such as interest and excitement, the amount of influence each person has on another in a social interaction, and unconscious back-and-forth interjections, while ignoring the words [23].
- Performing indoor user localization by measuring received signal strength and using triangulation algorithms that can achieve position estimation errors as low as 1.5 meters, which also allows for detection of people in close physical proximity [14].
- Communicating with Bluetooth enabled cell phones, PDAs, and other devices to study user behavior and detect people in close proximity.
- Capturing face-to-face interaction time using an IR sensor that can detect when two people wearing badges are facing each other within a 30°-cone and one meter distance. Choudhury [6] showed that it was possible to detect face-to-face conversations of more than one minute.

This represents a fundamental shift from earlier work in organizational behavior, since with this technology we are able to objectively quantify behavior at a level of detail unimaginable just a few years ago. In addition, we can examine radically different behavioral features than is possible using traditional observational and survey methods. Using

this data we hope to put the *social* back into organizational design and help people gain a better understanding of how their behaviors impact their performance and satisfaction at work.

To study how effective network structures differ in face-to-face networks, we deployed our Sociometric badge platform for a period of one month (20 working days) at a Chicago-area data server configuration firm that consisted of 28 employees, with 23 participating in the study. Each employee was instructed to wear a Sociometric badge every day from the moment they arrived at work until they left their office. In total we collected 1,900 hours of data, with a median of 80 hours per employee. All of these employees were male, and since this was a recently formed department, none had been employed for over a year. Still, there were five recognized experts, and in our analysis we controlled for skill level differences. Electronic communication was not extensively utilized in this firm for task-related communication, so we did not collect this data. Below, we explain the actual task structure for these employees, and in our analysis, we examine employee behavior at the task level rather than at the individual level. This allows for a much finer-grained analysis than would otherwise be possible, as well as uncovers some startling results.

Task Structure and Productivity Data

Salesmen in the field used an automated program to request a computer system configuration for a potential customer. These configurations are automatically assigned a difficulty (basic, complex, or advanced, in ascending order of difficulty) based on the configuration characteristics. Employees in the department are then assigned a configuration task in a first come first served fashion. This configuration task may require them to use a computer aided design (CAD) program in order to satisfy the customer’s needs. Finally, the employee submits the completed configuration as well as price back to the salesman, and the employee is placed at the back of the queue for task assignment. The exact start and end time of the task is logged, and the number of follow-ups that are required after the configuration is completed is also recorded in the database. We were able to obtain this data in addition to the badge data, although in our analysis, we only examined tasks where the employee was wearing the Sociometric badge for the entire task duration.

Measuring Cohesion

Network constraint C_i measures the degree to which an individual’s contacts are connected to each other. P_{ij} is the proportion of i ’s network time and energy invested in communicating with j . Network constraint can be used as proxy for measuring network cohesion [5], and network diversity is simply computed as $1-C_i$.

$$C_i = \sum_j \left(P_{ij} + \sum_q P_{iq} P_{qj} \right)^2, \quad q \neq i, j.$$

We expect face-to-face networks to require different network structures to transfer fundamentally different types of

knowledge when compared to email networks. Structurally diverse networks that use less rich media such as email are beneficial for obtaining diverse sources of information and consequently improving worker productivity [1]. Based on information richness theory and social network theories, cohesion (rather than diversity) in face-to-face networks should improve work performance as face-to-face communication is typically used to transfer more complex, embedded knowledge, and because network cohesion aids complex knowledge transfers. We therefore hypothesize that network cohesion is positively associated with work performance in face-to-face networks.

B. Analysis and Results

Network cohesion is positively correlated with work performance. Instead of reducing speed and productivity, as in email networks, a one-standard-deviation increase in network constraint in face-to-face networks is associated with a 9.5% increase in the speed of task completion, demonstrating that cohesive ties in a face-to-face network are more conducive to productivity than diverse ties. We suspect that the information transmitted in face-to-face networks is inherently different from that which is transferred in email networks. It appears that the advantages of using face-to-face communication to transmit complex knowledge are enhanced in cohesive networks. These results show that having a tight social group that can lend social support and enable trust to develop is extremely conducive to creating a more friendly and productive organizational environment.

IV. REALITY MINING AND PRIVACY

The preceding sections shows that we can draw rich inferences from peoples’ digital data—their relationships, their exposure to and likelihood of adopting social behaviors, and factors that impact and enhance their productivity. These data and computation tools hold the promise of socially aware applications and technologies. In the course of running these experiments, however, participants often raised important questions about their privacy and how this data would be used. Our results give us some hints as to how companies will make use of this kind of data in the future, so below we examine in detail the most common and pressing concerns.

A. ‘If I am an employee, does my company own my workplace behavior data?’

Technology used to monitor workplace interactions has the potential to increase general security and employee productivity, but there is also potential for disproportionate loss of workplace privacy. In general, the European Union has more stringent data privacy policies than the United States. The EU Directive on Privacy and Electronic Communications, which pertains specifically to public networks and public employees, claims that storage of individuals’ communication data is usually only permitted if users offer their explicit consent [10]. With regard to laws applicable to both the public and private sector, Article 8 of the European Convention on Human Rights maintains that “everyone has

the right to respect for his private and family life...and his correspondence,” which has mostly translated to courts upholding one’s right to privacy, even when one’s personal correspondence occurs through company-owned machines [16].

The United States, on the other hand, has a relatively loose policy concerning employee monitoring and has few laws that govern this area. There are few stipulations when it comes to storing data on company-owned servers because as long as there is no expectation of privacy, companies are permitted to access communications, such as e-mails, which are stored on their servers. In fact, as long as employees do not have a reasonable expectation of privacy—usually because employers have informed employees of possible monitoring—employers are allowed to monitor employees through forms such as phone, computer, and video surveillance [10].

Currently, the use of Sociometric badges in corporate settings seems to fall entirely within the scope of US law, as one can claim that the use of the badges is analogous to unconcealed video surveillance of a “publicly-accessible area,” since both noticeably gather data about people’s interactions, movement, and location [27]. After reviewing currently pertinent US and international laws concerning privacy in the workplace, it seems that the most appropriate approach to Sociometric badge data collection and storage would be to have third-party companies store badge data and implement the badge systems. In this case, the company utilizing the badges would not be in possession of the personally identifying raw badge data but could obtain certain network-level statistics about employees on both an aggregate and individual level. Assuming that this policy would be fully disclosed to employees, such measures would not only provide employers with useful metrics to help improve work culture and productivity, but they would also give employees a greater degree of privacy than the bare minimum required by US laws. We believe that, when using the badges, it would be best to follow guidelines established by the International Labour Organization, in which case badge data would only be lawfully collected and/or transferred with the informed consent of employees, workers would have access to their securely-stored personal badge data, and data would only be collected for “reasons directly relevant to employment” [27].

B. ‘If I am an individual, does my mobile operator or banking institution own my behavior data?’

Phone companies collect data about their users, called Consumer Proprietary Network Information, or CPNI. Although the phone companies own this data, the Federal Communications Commission, through the Telecommunications Act of 1996 Section 702, maintains fairly strict requirements on how the data is used. Companies are required to ensure the privacy of the data, and may only disclose the data to business affiliates who provide secondary services necessary for the telecommunications services being provided (all affiliates receiving data are required to keep the data as private as does the original company). They cannot

disclose the information to third parties for their own marketing purposes, but they are required, upon written request from the customer, to disclose the customer’s information to any party specified. The chief variation between phone companies’ privacy policies is whether the company adopts an opt-in or opt-out policy. Banks maintain similar privacy policies, as required by Regulation P: Privacy of Consumer Financial Information (12 CFR 216) of the Federal Reserve Board.

In addition to having the right not to have their data released, consumers also have the right to force telecommunication companies to release data on their behalf. If a third-party company, not licensed by the company, wants to use the data with the user’s permission, the user is only required to submit a written request designating the recipient (47 USC 222) and the telecommunication company is *required* to comply (e.g. if a user wants to use a third-party application for value-added analytics on his/her social data). With financial data, the laws empowering banks do not appear to *require* the release of information, although banks are *permitted* to release information at the direction of a consumer (12 CFR 216.15).

C. ‘If my data is anonymized, that means I’m safe, right?’

Many publicly released datasets rely on removing all personal identifiers from the data, in an attempt to *anonymize* the dataset so the participants cannot be identified, but this approach alone may not guarantee participant privacy. For pure social network data, Backstrom and colleagues [2] have proposed a family of attacks whereby it is possible to identify original participants with the help of embedded nodes. They suggest both passive and active forms of this attack, and identify 2400 edges in a 4.4 million-node network, by creating only 7 dummy nodes. Narayanan and Shmatikov [19] demonstrate a different method for passive de-anonymization by using a known auxiliary graph related to the *anonymized* dataset. From the legal perspective, the use of such anonymous data (i.e. personal identifiers removed) is not specified under the provisions of the US Electronic Communications Privacy Act, but was mentioned in the EU Directive on Privacy and Electronic Communications [Directive on Privacy].

D. ‘I’m not a participant or application user. Are you collecting any data about me?’

It is possible that Reality Mining applications may collect data unintentionally from non-participants and other third parties. Consider a simple example—a sensor that periodically scans and logs Bluetooth devices at a particular location. Individuals who are not application users may object that their unique Bluetooth identifiers are logged by the system. However the data being collected is public information and non-users are free to set their Bluetooth devices to non-discoverable mode (the default setting on most new phones and laptops, where Bluetooth communication is active but the unique identifier is not continually broadcast). Consider a legal precedent—the case of Smith v. Maryland [25]. Smith

established that, when making phone calls, there existed no legitimate expectation of privacy with respect to information such as the recipient or duration of a call. The Bluetooth argument is analogous, i.e. if you are broadcasting your Bluetooth identifier; you have no expectation of privacy with respect to your Bluetooth identity. Smith, though, applied to the internal records of the phone company contracted by the caller, and an argument can be made that there still exists a legitimate expectation of privacy with those companies not specifically contracted by the caller. Similarly, since the non-user has not contracted the application developer (e.g. by not using the application), the non-user has a legitimate expectation of privacy where the collection of his data by the application developers is concerned. Overall, in the Bluetooth case, the application developers would likely not be legally liable even if they collected Bluetooth ids of non-participants (under 652B, Intrusion Upon Seclusion) since the information is public and therefore carries no legitimate expectation of privacy. However as the data collected becomes progressively more private, the potential liability would increase.

V. CONCLUSION

In this paper, we describe two studies that show how data from mobile phones and Sociometric badges can help us understand human behavior. Mobile phone features can be used to accurately identify relationships between individuals. Similarly, these features can predict the future sharing of viral music within a social network. Similarly, badge features show that cohesive ties and higher face-to-face interactions are correlated with better task-level performance. These results illustrate the importance of behavioral data in both social and corporate settings.

There are many related privacy questions, and we have commented on a few of them. Within corporate settings, badge data today would be viewed as any other company property. Within social settings there are more limitations as far as service provider's ownership of the data, but not all companies make it easy for their customers to control sharing of their data, employing opt-out rather than opt-in policies. Furthermore, anonymization of data is often difficult, if not impossible, to achieve, and thus should not be considered a privacy safeguard. Finally, it is important to consider non-participants and third parties who may be affected during the data-collection process.

ACKNOWLEDGMENT

The authors would like to thank Lanthe Chronis and Daniel Olguin Olguin for their support with experiments.

REFERENCES

- [1] S. Aral, E. Brynjolfsson and M.W. Van Alstyne, "Productivity Effects of Information Diffusion in Networks," MIT Center for Digital Business, paper 234
- [2] L. Backstrom, C. Dwork, and J. Kleinberg. "Wherefore Art Thou R3579X? Anonymized Social Networks, Hidden Patterns, and Structural Steganography." *WWW Conference*, 2007.
- [3] W. E. Baker. *Achieving Success Through Social Capital: Tapping Hidden Resources in Your Personal and Business Networks*. Jossey-Bass, 2000.
- [4] H.R. Bernard, P. Killworth, D. Kronenfeld and L. Sailer, "The Problem of Informant Accuracy: The Validity of Retrospective Data", *Annual Reviews in Anthropology*, 1984.
- [5] R. S. Burt. *Structural Holes: The Social Structure of Competition*. Cambridge, MA USA, Harvard University Press, 1992.
- [6] T. Choudhury. "Sensing and Modeling Human Networks." Cambridge, MA USA, PhD Thesis, MIT Media Laboratory, 2004.
- [7] R. Cross and A. Parker. *The Hidden Power of Social Networks*. Boston, MA USA, Harvard Business School Publishing, 2004.
- [8] Directive on Privacy and Electronic Communications. *Official Journal*. L 201 , 31/07/2002 P. 0037 – 0047. EUR-Lex.
- [9] N. Eagle and A. Pentland, "Reality Mining: Sensing Complex Social Systems", *Personal and Ubiquitous Computing*, Vol 10, #4, 255-26.
- [10] "Employee Monitoring: Is There Privacy in the Workplace?" Fact Sheet 7: Workplace Privacy. Privacy Rights Clearinghouse, April 2009.
- [11] M. Gonzalez, C. Hidalgo and A.-L. Barabási. "Understanding Human Mobility Patterns." *Nature* 453, pp 779-782 2008.
- [12] Granovetter M. 1978. Threshold models of collective behavior. *American Journal of Sociology*, 83, 6, 1420-1443
- [13] F. Grippa, A. Zilli, R. Laubacher and P. Gloor. "E-mail may not reflect the social network." *Proceedings of the North American Association for Computational Social and Organizational Science Conference*. 2006.
- [14] Y. Gwon, R. Jain and T. Kawahara. "Robust indoor location estimation of stationary and mobile users." *Proceedings of the 23rd Annual Joint Conference of the IEEE Computer and Communications Societies*. 2004. 1032-1043.
- [15] Huckfeldt R, Sprague J, 1991. Discussant Effects on Vote Choice: Intimacy, Structure and Interdependence. *The Journal of Politics*, 53, 1, 122-158.
- [16] IBLs Editorial Department. "European Approach to Privacy in the Workplace." *Internet Business Law Services*. Retrieved 16 June 2008.
- [17] B. L. Kirkman, B. Rosen, P. Tesluk and C. B. Gibson. "The impact of team empowerment on virtual team performance: the moderating role of face-to-face interaction." *The Academy of Management Journal*, 2004: 175-192.
- [18] J. Leskovec, L. Adamic and B.A. Huberman. "The dynamics of viral marketing." *ACM Trans. Web.*, 1, 1, 2007
- [19] A. Narayanan and V. Shmatikov, "De-anonymizing Social Networks," *to appear at IEEE Security & Privacy '09*.
- [20] D. Olguin Olguin, B. N. Waber, T. Kim, A. Mohan, K. Ara, and A. Pentland. "Sensible Organizations: Technology and Methodology for Automatically Measuring Organizational Behavior." *IEEE Transactions on Systems, Man, and Cybernetics Part B*, 2009: 43-55.
- [21] J.-P. Onnela, J. Saramäki, J. Hyvönen, G. Szabó, D. Lazer, K. Kaski, J. Kertész and A.-L. Barabási. "Structure and Tie-strengths in Mobile Communication Networks." *PNAS* 104, 7332-7336.
- [22] A. Pentland. "Automatic mapping and modeling of human networks." *Physica A: Statistical Mechanics and its Applications*, 2006.
- [23] A. Pentland. "Socially Aware Computation and Communication." *IEEE Computer*, 2005: 33-40.
- [24] E.M. Rogers. *Diffusion of Innovations*. New York NY USA, Free Press, 1995.
- [25] Smith v. Maryland, <http://caselaw.lp.findlaw.com/scripts/getcase.pl?court=US&invol=735&vol=442>
- [26] B.N. Waber, D. Olguin Olguin, T. Kim and A. Pentland. "Understanding Organizational Behavior with Wearable Sensing Technology." *Academy of Management Annual Conference*. Anaheim, CA, USA, 2008.
- [27] "Workplace Privacy." EPIC Workplace Privacy Page. Electronic Privacy Information Center, 11 September 2008. Retrieved 14 July 2009.