

Collective Modeling of Human Social Behavior

Danny Wyatt

Dept. of Computer Science and Engineering
University of Washington
Seattle, WA 98195
danny@cs.washington.edu

Abstract

This paper presents my vision for collective behavior modeling: simultaneously modeling the behavior of all people within a group while taking into account the structure of their social ties. I believe that collective behavior modeling will provide new insights into the relationship between local individual behavior and global social structure while also improving the accuracy of predictions about human behavior. I summarize my work to date in the area, consider related work by others, and discuss the challenges and opportunities for future work.

Introduction

Human behavior is typically modeled only at the level of a single person (e.g. (Liao, Fox, and Kautz 2005; Hodges and Pollack 2007; Huynh, Fritz, and Schiele 2008)). When data is available for multiple people, each person is considered independent and predictions are made about each separately. Existing models of social behavior typically rise only to the level of the dyad (a single pair) (Pentland 2007) or small interacting group (Gibson 2005). On the few occasions that entire groups are modeled jointly (McCallum, Wang, and Corrada-Emmanuel 2007; Eagle, Pentland, and Lazer 2008), it is only to examine the immediate social context of each person and not the structure of interpersonal ties—the social network—that spans the entire group. In general, behavior modeling is concerned with only the local behavior around one person, and not the global social structure within which that behavior occurs.

Social network analysis, the study of network structure, has a long history (Radcliffe-Brown 1940; Granovetter 1973; Wellman and Berkowitz 1988) and has developed an arsenal of techniques for studying social structure (Wasserman and Faust 1994; Carrington, Scott, and Wasserman 2005). However, all of those techniques consider only the existence and organization of social ties and not the behavior that manifests itself along those ties. In general, social network analysis is concerned with only the global structure, and not the local behavior within that structure.

Recent advances in ubiquitous sensing and computing have made it possible to gather data about the simultaneous, real-world behavior of entire groups of people (Choudhury

and Pentland 2003; Wyatt, Choudhury, and Kautz 2007; Wren et al. 2007). Such data sets often capture, either directly or indirectly, interactions between people in the group. As such, they provide an entirely new view of *both* local individual behavior and global social structure.

Collective behavior modeling—simultaneously modeling the behavior of all people within a group while taking into account the structure of their social ties—can open up many exciting new research questions. For example, is there a relationship between a person’s behavior and how central she is to her network? Are there differences in behavior between people based on their network position or the configuration of ties in the network around them? Do clusters of behavior correspond to sub-groups within the social network, or to types of relationships between people? Can behavior be used to predict social position? Can behavior be used to predict which social ties will form or dissolve? Can network structure be used to predict how two people will interact?

Beyond exploring new sociological questions, the collective modeling of automatically measurable behavior data will also enable new applications that can take advantage of knowledge of a person’s social context or provide feedback about her social behavior. Collective behavior modeling may also improve the automated prediction and recognition of human behavior. Collective classification has improved results in many other domains by considering the conditional dependencies between the entities to be classified (Jensen, Neville, and Gallagher 2004; Sen et al. 2008). For collective behavior modeling, the social network can define dependencies between people’s behavior, and features of the social network may be used to improve prediction and recognition. Additionally, some of the statistical techniques developed for social network analysis may find new application to behavior modeling (collective or otherwise) and machine learning.

In the rest of the paper I discuss the work on collective behavior modeling that my colleagues and I have done to date, including the collection of a novel longitudinal data set capturing face-to-face interactions and our early results correlating local behavior with global structure. I consider related work by others and contrast it with my vision for collective behavior modeling. I outline plans for my continued work in collective behavior modeling: building unified models of local behavior and global structure, modeling change

in behavior and structure over time, predicting future behavior, and ensuring the feasibility of any new techniques. Finally, I close by considering the broader applicability of any methods developed for collective behavior modeling.

Current Work

My colleagues and I have collected a data set that captures the face-to-face conversations between two cohorts of incoming graduate students. Both cohorts were in the same department at a large research university. In the first cohort, 24 of 27 eligible subjects participated. In the second, 17 of 29 possible subjects participated.

Each subject wore a sensing device containing 8 different sensors useful for detecting conversations, activities, and environmental context. Data was collected during working hours for one week each month over the 9 month course of an academic year. To collect data in an ethical (and legal) manner, no raw audio was ever recorded. Only a set of privacy-sensitive features were saved. At the end of each data collection week, the subjects submitted surveys reporting on their interactions with other participants during the previous month. A complete description of the data for one of the cohorts is in (Wyatt, Choudhury, and Kautz 2007).

This data set is novel in several respects. First, it directly captures real-world face-to-face conversations, which remain people’s primary mode of social interaction (Baym, Zhang, and Lin 2004). While, there have been a few earlier efforts towards the direct recording of face-to-face interaction, those required human observers and manual coding (Bernard, Killworth, and Sailer 1980; Freeman, Freeman, and Michaelson 1988)—techniques that can only be applied to small study populations over brief observation periods.

A second novel aspect of our data is that it is longitudinal. It is difficult to observe real-world interactions at even a single point in time; multiple observations at many different time points are clearly even more difficult.

Which brings up a third novel aspect of our data: it is automatically collected and processed. Automated recording and processing not only increases the scale—both in number of subjects and length of observation period—at which interactions can be studied, it also makes possible applications that have access to real-time information about a group’s social network.

Conversation Detection and Speaker Segmentation In earlier work my advisors and I have developed techniques for determining from our privacy-sensitive features who is colocated with whom, who is in conversation with whom, and who speaks when in a conversation. All of those techniques involve lower level probabilistic models whose outputs are fed into each other to ultimately produce high level inferences about conversations. Those techniques are explained fully in (Wyatt, Choudhury, and Bilmes 2007).

Ultimately, the high level inference results in a rich corpus of data about interactions and behavior. We can infer which subjects are physically located together, who speaks with whom, when and for how long. Additional features—pitch and energy—capture non-linguistic properties of how people speak. Altogether, the data contain information shown

Table 1: Correlation between change in speech and centrality

Rate		Pitch		Turn Length		Turn Frequency	
r	p	r	p	r	p	r	p
.307	.0003	.228	.0075	.164	.0558	-.0413	.6334

to be useful for inferring emotion (Liscombe, Venditti, and Hirschberg 2003), interest level (Gatica-Perez et al. 2005), and mental state (Hurlburt, Koch, and Heavey 2002).

Additional Behavior The other sensors in our data can be used to infer the wearer’s physical activity (e.g. walking, sitting, standing, etc.) and whether she is indoors or outside (Lester et al. 2005). That can allow us to see whether one person walks to another’s location, or whether people move (together or separately) to a new location. We can also use physical activity as a feature for conversation analysis: are the conversants sitting still or standing and moving about?

Early Results

This section presents two early results based on preliminary analyses of the data from the first cohort (24 subjects). For both results, only the 6 consecutive months with the most time recorded (3,021 hours) were used.

Conversational Behavior and Network Ties In (Wyatt et al. 2008) we present an early example of the new kind of question that we can answer about the correspondence between social structure and individual behavior: do individuals change their speaking styles more when interacting with people who are more central to the network?

Speaking style is taken to be one of four measurable features: pitch, rate, turn length, and turn frequency. For person i , we compute 3 quantities. (1) $b_{i \setminus j}$: the mean of feature b for i when i speaks with *everyone except* person j , (2) $b_{i \rightarrow j}$: the mean of b for i when i speaks *only* with j , and (3) s_i : the standard deviation of b for i , regardless of conversation partner. Let $d_{ij} \triangleq |b_{i \setminus j} - b_{i \rightarrow j}| / s_i$ be the amount that i ’s behavior changes when in conversation with j . The mean of d_{ki} for all k who speak to i is i ’s mean *incoming* change. The higher this incoming mean, the more people change their speaking style when with person i .

How central each person is to the network is computed using a variant of closeness centrality (Wasserman and Faust 1994), modified to take advantage of the unique continuous measure of social behavior in our data. Let c_{ij} be the proportion of time that i and j spend in conversation. The “length” of an edge between i and j is defined as $1 - c_{ij}$, if $c_{ij} > 0$. If i and j do not converse at all, their edge is null and its length is undefined. Closeness centrality is defined as the multiplicative inverse of the mean path length (via the shortest path) from a person to all others.

We found that there is indeed a positive correlation between a person’s mean incoming change and her closeness centrality for all speech features except turn frequency (Table 1). This supports an earlier finding by one of my advisors that there is a positive correlation between the amount of influence a person exerts on the turn-taking pattern of a conversation and his centrality (Choudhury and Basu 2004).




Figure 1: Closeness centrality and change in conversational behavior. Distance from the center indicates decreasing centrality. Darker color and larger size indicate higher incoming change. Lines connecting people indicate that they spent time in one-on-one conversation.

Figure 1 is a visualization of this comparison for speaking rate using the data from one observation week. The larger, darker nodes have higher incoming change and are found closer to the center of the network, while the smaller, lighter nodes are found out towards the edges.

Recovering Latent Networks from Conversations Traditional social network analysis assumes that social ties are observable and that a social network is directly available for study. When the data to be analyzed represents real-world behavior, however, it becomes necessary to define “some means of abstracting from these empirical acts to relationships or ties” (Marsden 1990). Typically, researchers define simple thresholds or heuristics to discard observations that are believed *a priori* to be noise. The remaining observations are then considered a direct observation of the “true” network (e.g. (Palla, Barabási, and Vicsek 2006; Kossinets and Watts 2006)) and this surviving network is used as data for subsequent analysis. Such methods do not propagate any uncertainty due to noisy observations into later analyses.

Another approach is to consider the social network a hidden state and the measurable behavior a noisy observation of that hidden state. In (Wyatt, Choudhury, and Bilmes 2008) we combine exponential random graph models (ERGMs), a state-of-the-art method from social network analysis, with the traditional statistical machine learning technique of using latent variables to model hidden state.

ERGMs (also called p^* models) model a social network as a realization of a set of random variables, one variable for each potential edge in the network (Frank and Strauss 1986; Wasserman and Pattison 1996; Robins et al. 2007). Given an observed network, ERGMs estimate the parameters of a

model that describes the joint distribution of the edge variables. The distribution takes the form of a log-linear combination of features and weights: $p(\mathbf{Y} = \mathbf{y}) = \frac{1}{Z_\theta} e^{\theta^\top \phi(\mathbf{y})}$, where $\mathbf{Y}$ are the variables representing edges in the graph, ϕ are feature functions, θ is a vector of weights to be learned, and Z_θ is the usual normalizing constant. The features are deterministic functions of the variables and they define conditional independence assumptions between variables. Simple features are typically counts of occurrences of specific subgraphs within the network, e.g. the number of triangles, or even just the number of edges.

More complex feature sets involve histograms of empirical distributions of statistics (e.g. degree) of the network. To avoid overfitting, the weights for these histogram features are constrained to follow pre-defined functions (Hunter 2007). The constrained model has fewer parameters while also capturing intuitions about the smoothness that should exist between weights on features that have an intrinsic relationship. For example, the weights on bins in the degree histogram increase according to degree, but at a geometrically decreasing rate. That models the intuitive notion of a diminishing rate of return for adding new ties.

For our model, we added a new set of variables corresponding to observable behavior for each pair of subjects in our population—specifically the amount of time that each pair spends in conversation. These observable variables are each connected to a latent variable that models whether or not a social tie exists between that pair. By marginalizing out the network—treating it as a distribution over networks and not one single, observable network—we carry any uncertainty inherent in the observations through to the subsequent analysis of the network.

We used the sociological features most common in ERGM-based analyses together with a new set of features tying observable behavior to latent network state. Our new features also make use of non-linear parameter constraints to model the same “diminishing returns” intuition about time spent in conversation. The sociological features represent the global properties of the network while the observation features represent the local properties of individuals. Importantly, the sociological features span the entire network so that no pair is considered marginally independent of any other pair. Conditional dependencies between pairs allow behavioral information to spread through the model so that the resulting parameters simultaneously balance information about all pairs’ behavior with a consistent (according to the sociological features) model of how the pairs fit together into a larger network.

The resulting model recovered networks that had better agreement with the networks reported through surveys than the basic network of conversations alone. More importantly, the model was able to learn—in a fully unsupervised manner—properties of the connection between conversational behavior and the structure of social ties within the group. Figure 2 shows two conditional probabilities learned by the model. The solid line is the probability that a pair will spend a certain amount of time in conversation given that a social tie exists between them. The dashed line




Figure 2: Conditional probabilities of time in conversation given existence or non-existence of social tie.

is the probability that a pair will spend time in conversation given that no tie exists between them. For this population, the threshold for “socially significant” time in conversation appears to be about 6 minutes. Additionally, it appears that the social utility of time in conversation begins to diminish after about 12 minutes. Note that the entries along the x-axis actually represent discrete categories in a multinomial, but the conditional probabilities follow smooth curves because of the parameter constraints.

Related Work

The work most similar to ours is that of the MIT Reality Mining project (Eagle, Pentland, and Lazer 2008) along with similar work from the same group (Olguin et al. 2009). The Reality Mining data contains measurements of the physical proximities (via Bluetooth radio signals) and cell phone calls within a group of 100 students over the course of one year. Their newer data measures the face-to-face, line-of-sight proximity between 22 bank employees, as well as all of their email interaction.

(Eagle, Pentland, and Lazer 2008) summarizes 3 approaches to using the Reality Mining data to predict a subject’s self-reported ties, satisfaction, and recall of proximities based on her sensed proximities. All of the predictive models involve features at either the level of the dyad or the immediate social context around one subject. These results might be strengthened by considering more complex properties of the network. For example, whether a person’s friends are also friends with each other and whether and how often such a group of mutual friends comes together may affect that person’s satisfaction or recall. One analysis of their bank employee data does consider global network structure by examining the relationship between network centrality and reported satisfaction.

None of these analyses address the difficulty of using location data as a proxy for true interaction. The Reality Mining subjects are students who attend classes and lectures and work in neighboring offices. That they are physically close does not mean they are actually interacting. (And since their proximity is measured from radio signals, subjects who are sensed to be close to one another may not even be in the same room!) There may be large differences between a network built from proximity data and one built from data about real interactions. Our data set captures both conversations and colocation, so we can study the difference between analyses that use only proximity data and those that use actual

interaction data.

Another recent similar effort is that of (Connolly, Burns, and Bui 2008). They show that events that may be social interactions—attending the same meeting, visiting someone’s office, walking together—can be extracted from the motion sensor data of (Wren et al. 2007). As with the Reality Mining data, the “interactions” derived from motion sensor data are very ambiguous and any further analysis will need to account for that uncertainty.

In the social science literature, (Gibson 2005) studies the relationship between turn-taking patterns in meetings and the reported social ties between participants (e.g. supervisor/subordinate, peer coworkers, friends). His data had to be manually coded by two human observers (in real-time, no less)—an approach that obviously cannot scale to large groups or longitudinal studies. Furthermore, his analysis only considers the relationships between dyads separately and does not extend to consider the structure of relationships between all of the people in the study.

The work on topic models in email (McCallum, Wang, and Corrada-Emmanuel 2007) comes closest to jointly modeling the entire behavior of a group, but obviously the behavior examined is the generation of written text and not face-to-face interaction. Social ties are not explicitly modeled, but behavior is considered between interacting pairs and models are fit to all pairs simultaneously. However, while the social ties are exploited, potentially useful features of the global structure like reciprocity or transitivity are not.

Plans for Future Work

Our early result relating change in conversational style to centrality used a disconnected chain of very different statistical methods. The main goal of the work in my thesis will be to formulate a unified statistical model that ties global social structure to local individual behavior. Ideally, the model will be capable of both explaining and predicting human social behavior. That is, it will be capable of learning parameter values that can be interpreted in sociologically meaningful ways, and it will be capable of using those learned parameters to predict behavior in out-of-sample data.

Constructing such a unified model can be thought of in two equivalent ways. From the perspective of social network analysis, it can be seen as extending ERGMs “downwards” so that they include behavior data. From the perspective of statistical behavior modeling, it can be seen as extending behavior models “upwards” so that they can take advantage of the dependencies between people defined by the structure of their social network. As such, a unified model will hopefully take advantage of techniques developed separately for either behavior modeling or social network analysis, as well as producing new techniques potentially useful to each of them independently.

The work I propose for my thesis can be roughly separated into four intersecting areas of exploration: (1) methods for tying concrete behavior to abstract social relationships, (2) modeling changes in behavior and social structure over time, (3) predicting future behavior based on collective observations, and (4) ensuring the tractability of all the techniques developed.

From Behavior to Relationships

As mentioned previously, behavior is observable, social ties are not. Our earlier work learned a probabilistic relation between time spent in conversation and the existence of a social tie. Future models must incorporate more behavior data. They should consider how frequently a pair interacts, how they speak and take turns during their interaction, who moves to visit whom, and even the time of day and day of the week of their interactions. All of these behavioral features must be associated with an abstract notion of social tie, and the representation of that tie may need to become more complex. Ties may need to be modeled as categorical (friends, coworkers) or continuous-valued (to indicate strength of tie) variables—or possibly without any explicit variable at all (discussed below).

Additionally, while social networks are modeled as structures of dyads, many interactions occur between more than two people. To model multiparty connections, traditional social network analysis has used bipartite graphs where one set of nodes represents people and the other represents interactions. ERGMs have been extended to handle such bipartite graphs (Wang 2006), however the features used no longer model a structure of abstract social ties but rather the structure linking people to interactions. It may be beneficial to continue to model a social structure composed of dyads while also modeling connections between dyads and their conversations.

Managing Uncertainty Another concern when tying behavioral data to social structure is that most behavioral data does not directly represent true interactions but rather some proxy that may indicate an unobserved interaction. Sometimes, individual people are not identifiable (e.g. in motion sensor data) and sometimes the data does not represent “social reality” (e.g. proximity readings through walls). The data that my group has collected captures information about interactions only indirectly: the interactions must be inferred from raw sensor output. A low-level model (explicit or implicit) that infers interactions from sensor data should preserve any uncertainty found in the original observations and propagate that uncertainty to higher level analyses.

An additional source of uncertainty comes from the spatially distributed nature of the sensors. Observing an entire group’s behavior over time means having to observe multiple events occurring in separate locations. That means that multiple points of observation, possibly with varying measurement error rates, must be aggregated. Data could even be missing from some sensors but present from others. (Indeed, there is never a time in our data when all subjects are recording simultaneously.) Uncertainty due to uneven measurement error and missing data must also be propagated to higher levels of analysis.

Temporal Models

Sensor data is inherently temporal and future models should make use of that temporality. Temporal extensions to ERGMs exist (Pattison and Robbins 2001; Guo et al. 2007), generally treating the network as a process that evolves in discrete time steps with a Markov assumption of conditional

independence between non-successive time steps.

Connecting such models to sensor data will require multi-rate models that allow the low-level behavioral processes (e.g. speaker turns) to evolve at a faster rate than the high-level social network process. Multirate models that aggregate fast evolving low-level processes into slowly evolving high-level processes have been used in behavior modeling before (Liao, Fox, and Kautz 2004). However, if the learned parameters are to provide sociological insight, it may be necessary to allow the high-level process to be time-inhomogeneous—something not done in existing behavior models.

An alternative model would allow the higher network level to evolve at the same rate as the lower sensor level while constraining high level changes to follow a smooth, slowly changing function. Allowing fast but small changes at the network level would be useful for capturing the complex nature of social ties. For example, there is no one moment when a tie should be “switched off” between friends who have fallen out of touch. Instead, the intensity of their relationship should fade with time.

Predicting Behavior

Models that have sufficient interpretability may use a very different set of features from those that produce accurate predictions. It may even be the case that prediction is possible without explicitly representing social structure (e.g. with specific latent variables) at all. Understanding the difference between models that predict well and models that are interpretable will hopefully also provide new insight into the connection between behavior and social structure.

Tractability

Larger models entail increased complexity for learning and inference. Treating the social network as a hidden state makes the problem of parameter estimation non-convex. Stochastic gradient ascent seemed to work in our early experiments, but those involved separated data points (one per observation week) that could be used in traditional stochastic gradient methods. A temporal model that ties observations together will effectively have only one large observed data point, and the typical stochastic gradient methods will no longer apply.

It may be beneficial to model only observable behavior without explicitly modeling unobservable social structure. A network constructed of separate behavioral measurements would require techniques for analyzing valued networks (those with non-binary edges). That will require new features that model behavior but are based on the same sociological intuition behind the features currently in use.

Additionally, the amount of data available at the low-level is enormous and connecting it to a large, unified model will not produce a model whose parameters can be tractably learned. The spatially distributed nature of the process may allow some parts of the model to be decomposed and managed in parallel. The trade-offs between accuracy and tractability involved in various model decompositions or approximations will need to be examined.

Broader Applicability

Hopefully the techniques developed and lessons learned while exploring the above 4 areas will have uses beyond collective behavior modeling of sensor-based behavior data. Any new models could also be applied to large corpora of virtual interaction data like email or instant messaging logs, or community website (e.g. Facebook) traffic. The features tying individual behavior to the global network may need to be changed, but hopefully the general class of models as well as the network-level sociological features will still be useful for analysis and prediction in other data sets.

Additionally, ERGMs can be seen as a specific application of techniques from the domain of statistical relational learning (SRL). ERGMs restricted to simple features could be implemented using existing techniques such as relational Markov networks (Taskar, Abbeel, and Koller 2002) or Markov logic (Domingos et al. 2008). The more complex features and non-linear parameter constraints are techniques currently unused in SRL. I hope that techniques developed for applying such constrained models to collective behavior modeling will have broader applicability to existing SRL methods such as hybrid Markov logic networks (Wang and Domingos 2008).

References

- Baym, N.; Zhang, Y. B.; and Lin, M. C. 2004. Social interactions across media: Interpersonal communication on the internet, telephone and face-to-face. *New Media and Society* 6.
- Bernard, H. R.; Killworth, P. D.; and Sailer, L. 1980. Informant accuracy in social network data IV: a comparison of clique-level structure in behavioral and cognitive network data. *Social Networks* 2(3).
- Carrington, P. J.; Scott, J.; and Wasserman, S., eds. 2005. *Models and Methods in Social Network Analysis*. Cambridge UP.
- Choudhury, T., and Basu, S. 2004. Modeling conversational dynamics as a mixed memory markov process. In *Proc. of NIPS*.
- Choudhury, T., and Pentland, A. S. 2003. Sensing and modeling human networks using the sociometer. In *Proc. of the Int'l Conference on Wearable Computing*.
- Connolly, C. I.; Burns, J. B.; and Bui, H. H. 2008. Recovering social networks from massive track datasets. In *Proc. IEEE Workshop on Appl. of Comp. Vision*.
- Domingos, P.; Kok, S.; Lowd, D.; Poon, H.; Richardson, M.; and Singla, P. 2008. Markov logic. In Raedt, L. D.; Frasconi, P.; Kersting, K.; and Muggleton, S., eds., *Probabilistic Inductive Logic Programming*. Springer.
- Eagle, N.; Pentland, A. S.; and Lazer, D. 2008. Inferring social network structure using mobile phone data. (*in submission*).
- Frank, O., and Strauss, D. 1986. Markov graphs. *J. Am. Statistical Association* 81.
- Freeman, L. C.; Freeman, S. C.; and Michaelson, A. G. 1988. On human social intelligence. *J. Soc. and Biol. Structures* 11.
- Gatica-Perez, D.; McCowan, I.; Zhang, D.; and Bengio, S. 2005. Detecting group interest-level in meetings. In *Proc. of ICASSP*.
- Gibson, D. R. 2005. Taking turns and talking ties: Networks and conversational interaction. *Am. J. Sociol* 110(6).
- Granovetter, M. 1973. The strength of weak ties. *Am. J. Sociol* 78(6).
- Guo, F.; Hanneke, S.; Fu, W.; and Xing, E. C. 2007. Recovering temporally rewiring networks: A model-based approach. In *Proc. of ICML*.
- Hodges, M. R., and Pollack, M. E. 2007. An 'object-use fingerprint': The use of electronic sensors for human identification. In *Proc. of UbiComp*.
- Hunter, D. 2007. Curved exponential family models for social networks. *Social Networks* 29(2).
- Hurlburt, R.; Koch, M.; and Heavey, C. 2002. Descriptive experience sampling demonstrates the connection of thinking to externally observable behavior. *Cognitive Therapy and Research* 26(1).
- Huynh, T.; Fritz, M.; and Schiele, B. 2008. Discovery of activity patterns using topic models. In *Proc. of UbiComp*.
- Jensen, D.; Neville, J.; and Gallagher, B. 2004. Why collective inference improves relational classification. In *Proc. of ACM SIGKDD*.
- Kossinets, G., and Watts, D. J. 2006. Empirical analysis of an evolving social network. *Science* 311.
- Lester, J.; Choudhury, T.; Kern, N.; Borriello, G.; and Hannaford, B. 2005. A hybrid discriminative-generative approach for modeling human activities. In *Proc. of IJCAI*.
- Liao, L.; Fox, D.; and Kautz, H. 2004. Learning and inferring transportation routines. In *Proc. of AAAI*.
- Liao, L.; Fox, D.; and Kautz, H. 2005. Location-based activity recognition. In *Proc. of NIPS*.
- Liscombe, J.; Venditti, J.; and Hirschberg, J. 2003. Classifying subject ratings of emotional speech using acoustic features. In *Proc. of Eurospeech*.
- Marsden, P. V. 1990. Network and data measurement. *An. Review of Sociol.* 16.
- McCallum, A.; Wang, X.; and Corrada-Emmanuel, A. 2007. Topic and role discovery in social networks with experiments on enron and academic email. *JAIR*.
- Olguin, D. O.; Waber, B. N.; Kim, T.; Mohan, A.; Ara, K.; and Pentland, A. S. 2009. Sensible organizations: Technology and methodology for automatically measuring organizational behavior. *IEEE Trans. on Sys., Man, and Cybernetics, Part B* 39(1).
- Palla, G.; Barabási, A.-L.; and Vicsek, T. 2006. Quantifying social group evolution. *Nature* 446.
- Pattison, P., and Robbins, G. 2001. Random graph models for temporal processes in social networks. *J. Math. Sociol.* 25.
- Pentland, A. S. 2007. Automatic mapping and modeling of human networks. *PhysicaA*.
- Radcliffe-Brown, A. R. 1940. On social structure. *J. Royal Anthropological Institute of GB and Ireland* 70(1).
- Robins, G.; Snijders, T.; Wang, P.; Handcock, M.; and Pattison, P. 2007. Recent developments in exponential random graph (p*) models for social networks. *Social Networks* 29(2).
- Sen, P.; Namata, G. M.; Bilgic, M.; Getoor, L.; Gallagher, B.; and Eliassi-Rad, T. 2008. Collective classification in network data. *AI Magazine* 29(3).
- Taskar, B.; Abbeel, P.; and Koller, D. 2002. Discriminative probabilistic models for relational data. In *Proc. of UAI*.
- Wang, J., and Domingos, P. 2008. Hybrid markov logic networks. In *Proc. of AAAI*.
- Wang, P. 2006. Exponential random graph (p*) models for affiliation networks. Master's thesis, University of Melbourne.
- Wasserman, S., and Faust, K. 1994. *Social Network Analysis*. Cambridge UP.
- Wasserman, S., and Pattison, P. 1996. Logit models and logistic regression for social networks: 1. an introduction to markov graphs and (p*). *Psychometrika* 61.
- Wellman, B., and Berkowitz, S. D., eds. 1988. *Social Structures: A Network Approach*. Cambridge UP.
- Wren, C. R.; Ivanov, Y. A.; Leigh, D.; and Westhues, J. 2007. The merl motion detector dataset. In *MD '07: Proc. of the 2007 workshop on Massive datasets*.
- Wyatt, D.; Choudhury, T.; Bilmes, J.; and Kitts, J. A. 2008. Towards the automated social analysis of situated speech data. In *Proc. of UbiComp*.
- Wyatt, D.; Choudhury, T.; and Bilmes, J. 2007. Conversation detection and speaker segmentation in privacy sensitive situated speech data. In *Proc. of Interspeech*.
- Wyatt, D.; Choudhury, T.; and Bilmes, J. 2008. Learning hidden curved exponential random graph models to infer face-to-face interaction networks from situated speech data. In *Proc. of AAAI*.
- Wyatt, D.; Choudhury, T.; and Kautz, H. 2007. Capturing spontaneous conversation and social dynamics: A privacy-sensitive data collection effort. In *Proc. of ICASSP*.