

Automatic Mapping and Modeling of Human Networks

Alex (Sandy) Pentland

Room E15-387,

Massachusetts Institute of Technology

Abstract: Mobile telephones, company ID badges, and similar common devices form a sensor network which can be used to map human activity, and especially human interactions. The most informative sensor data seem to be measurements of person-to-person proximity, and statistics of vocalization and body movement measurements. Using this data to model individual behavior as a stochastic process allows prediction of future activity, with the greatest predictive power obtained by modeling the interactions between individual processes. Experiments show that between 40 and 95% of the variance in human behavior may be explained by such models.

1. Introduction

A series of studies on office interactions discovered that 35-80% of work time is spent in spoken conversation, 14-93% of work time is spent in opportunistic communication, and 7-82% of work time is spent in meetings [1]. Senior managers represent the high end of these scales. Given the importance of such communications, it is notable that the majority of adults already carry a microphone and location sensor in the form of a mobile phone, and that these sensors are packaged with computational horsepower similar to that found in desktop computers. This emerging foundation of wearable sensing and processing power has allowed us to begin to automatically map

and model how different groups within social or business institutions connect.

We have been particularly concerned with our ability to automatically infer properties of human networks that affect propagation of information:

- Location Context: work, home, etc.
- Social Context: with friends, co-workers, boss, family, etc.
- Social Interaction: are you displaying interest, boredom, friendliness, determination, etc.

By taking a statistical, machine-learning approach applied to the users' behavior and physical situation, we have been able to show that it is possible to obtain solid, dynamic estimates of the users' group membership and the character of their social relationships: e.g., who we work for versus those who work for us, when we are interested versus when we are bored. By characterizing these patterns of behavior using statistical learning methods, we can then examine the users' current behavior to classify relationships as workgroup, friend, interesting, and so forth.

The key to automatic inference of information network parameters is the recognition that humans are not general-purpose equipotent reasoning agents, but rather are creatures with a long evolutionary history that continues to shape our behavior and interactions with others. This shaping of behavior is particularly visible in social relationships and our attitudes toward them:

we act differently when interacting with friends versus strangers, and when we are interested versus when we are bored.

Some of these categories can be inferred using standard methods such as surveys, however these standard methods often suffer from subjectivity and memory effects, and their infrequency means that they are prone to becoming out-of-date. Even when information from standard methods is available, we would still like to use automatic methods to validate or even correct the standard information sources.

In this paper, we present statistical learning methods that use wearable sensor data to make reliable estimates about a user's interaction state (e.g. who was talking to whom, how long did the conversation last, etc.). We then use these results to characterize the connections that exist in groups of people.

Automatic mapping can be much cheaper and more reliable than human delivered questionnaires. For instance, in one of our studies we found that our automatic methods had an accuracy of 87.5% for detecting conversations with durations of one minute or more. In contrast, a traditional survey of the same subjects produced only 54% agreement between subjects (where both subjects acknowledged having the conversation) and only 29% agreement in the number of conversations [2, 3].

Automatic discovery and characterization of face-to-face communication networks automatically will also allow researchers to gather interaction data from larger groups of people. This can potentially remove two of the current limitations in the analysis of human networks: the number of people that can be surveyed, and the frequency with which they can surveyed.

Automatic mapping of human networks will never be perfect, although already seems superior to previous methods in some regards. We can also vary the confidence thresholds of the system, making the system more or less cautious about particular types of mistakes. In addition, the models provided by automatic mapping can suggest when traditional survey methods would be most useful, resulting in a semi-automatic capability that can have very high accuracy and a relatively cost.

2. Socioscopes

Our approach to mapping and modeling human networks is to adopt the conceptual framework used in biological observation, such as is used to study apes in natural surroundings or in natural experiments such as twin studies, but replacing expensive and unreliable human observations with automated, computer-mediated observations. We imagined an advanced 'socioscope' that can accurately and continuously track the behavior of hundreds of humans at a time, recording even the finest scale behaviors with near perfect accuracy.

My students and I have built an approximation of this imaginary socioscope, using mobile telephones, electronic badges, and PDAs [4-8]. My collaborators and I have used this socioscope to track the behavior of 94 people in two divisions of MIT, the business school and the Media Laboratory, a group of 110 international researchers attending meetings at MIT, and certain other smaller groups in the wider Boston community. The subjects were typically between 23 and 39 years of age, with the business school students almost a decade older than the Media Lab students. Subject groups were typically 2/3 male and 1/3 female, and approximately half were raised in America.

The socioscope consists of three main parts. The first part consists of ‘smart’ phones programmed to keep track of their owners’ location and their proximity to other people, by sensing cell tower and Bluetooth IDs. This has provided us with approximately 330,000 hours of data covering the behavior of 94 people, a total of about 35 years of interaction data, as described in Eagle and Pentland [5].

The second part of the socioscope consists of electronic badges that record the wearers’ location (with 2 meters typical accuracy), ambient audio, and upper body movement via a 2-D accelerometer, as described in Gips and Pentland [6]. This badge platform provides more fine-grained data than the smart phone platform. We have used this platform to obtain data from the more than 110 adults that regularly attend the biannual Media Lab open houses, in which attendees walk around the Media Lab building to examine demonstrations and converse with each other during a four-hour period. The attendees have been approximately 1/3 from Asia, 1/2 from North America, and 1/6 from Europe.

The third part of the socioscope consists of a microphone and software that is used to extract audio ‘signals’ from individuals, specifically, the exact timing of their vocalizations and the amount of modulation (in both pitch and amplitude) of those vocalizations. This part of the socioscope can be used with audio data from the smart phone, audio from the badge, or audio from body-worn microphones during semi-structured interactions such as speed dating, focus group interviews, or negotiations [7,8].

Together these sensor platforms allow us to observe gross behavior (location, proximity) continuously over months, to observe for fine-grained behavior (location, proximity, body motion) over one-day

periods, and to analyze vocalization statistics with an accuracy of tenths of seconds. These behavioral data are then subject to four main types of analysis: characterization of individual and group distribution and variability (typically using an Eigenvector or principal components analysis), conditional probability relationships between individual behaviors (known as ‘influence modeling’), accuracy with which behavior can be predicted (with equal type I and II error rates), and finally comparison of these behavioral measures to standard human network parameters.

3. Reality Mining

A critical requirement for automatic mapping and modeling of human networks is to learn and later categorize user behavior as quickly as possible. This is because the speed with which we can establish network parameters determines how accurately we can capture the dynamics of those networks.

3.1. Eigenbehavior Modeling

Eigenvector analysis, commonly known as principal components analysis, is the optimal linear method for obtaining a low-dimensional approximation to a signal such as observations of user behavior. *Eigenbehaviors* thus provides us with an efficient method of learning and classifying user behavior, as described in Eagle and Pentland [9].

Calculation of Eigenbehaviors begins by measuring person I ’s behavior (for instance, the time sequence of their location). This is illustrated in Figure 1 as $\Gamma_1(x,y)$, a two-dimensional D by 24 array of location information, where D is the total number of days person I has been in the study. Because of the structure in most people’s lives, Γ_1 can be described by a relatively low dimensional ‘behavior space’, which is spanned by their Eigenbehaviors.

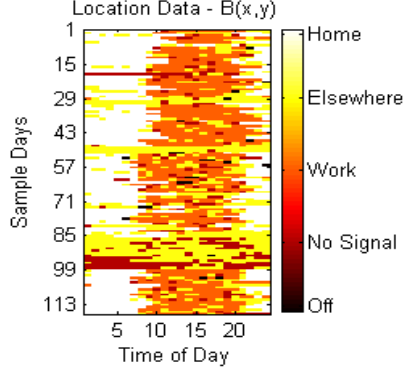


Figure 1. A sample of a person's location in a time sequence.

Given behaviors $\Gamma_1, \Gamma_2, \Gamma_3 \dots \Gamma_M$, for a group of M people, the average behavior of the group can be defined by $\Psi = \frac{1}{M} \sum_{n=1}^M \Gamma_n$. A set of M vectors, Φ_i , are defined to be the deviation of an individual's behavior from the mean: $\Phi_i = \Gamma_i - \Psi$. Principle components analysis is subsequently performed on these vectors generating a set M orthonormal vectors, u_n , which best describes the distribution of the set of behavior data when linearly combined with their respective scalar values, λ_n . These vectors and their corresponding scalars are the Eigenvector and Eigenvalues of the covariance matrix of Φ , the set's deviation from the mean.

$$C = \frac{1}{M} \sum_{n=1}^M \Phi_n \Phi_n^T = AA^T$$

where the matrix $A = [\Phi_1, \Phi_2, \Phi_3, \dots, \Phi_M]$. The Eigenbehaviors can then be ranked by the total amount of variance in the data for which they account, essentially those with the largest associated Eigenvalues.

3.2. Human Eigenbehaviors

As might be expected, the main daily pattern that we observed is that of subjects leaving their sleeping place to spend time in a small set of locations during the central daylight hours, then occasionally breaking into small clusters to move to one of a few other buildings during the early night hours

and weekends, and then back their sleeping place. Over 85% of the variance in the behavior of low entropy subjects can be accounted for by the mean vector alone.

For typical individuals the top three Eigenbehavior components account for up to 96% of the variance in their behavior. These three components could be thought of as the weekend pattern, the working late pattern, and the socializing pattern. Even though we are considering mostly young people without a regular job or family, it seems that there is very limited variability in human behavior. This means that one can, for instance, observe a person's behavior in the morning, and from these observations accurately predict their behavior for the rest of the day [9].

The ability to accurately characterize peoples' behavior with a low-dimensional model means that we can automatically classify the users' location context with high accuracy. If we also allow the system to request that the user label locations that the system thinks are novel or are unusually sequenced, then we can achieve very high accuracies with very limited user input.

3.3. Learning Influence

The previous data illustrate the stereotypical patterns and limited variability we observe in individual subjects. Next let us ask what behavioral structure we observe between subjects. We thus move from a static analysis of behavior to a dynamic analysis.

Conditional probability relationships between subjects, which we refer to as influence, allow us to predict the behavior of a subject from the other subjects' data [15,16]. For instance, if Joe shows up at a meeting whenever Fred does, then observing Fred's attendance allows accurate prediction of Joe's impending proximity. In our cell phone proximity data there were

two main sub-networks of influence relations, one during the day and the other in the evening, both with similar network prediction accuracy. Overall, influence between subjects allowed 95% of the variance in personal proximity data to be accounted for by the surrounding network of proximity data, as described in Dong and Pentland [11].

Again, a critical requirement for automatic mapping and modeling of human networks is the ability to learn and later categorize user behavior from relatively few observations, in order that we can accurately capture the dynamics of the network. The requirement for a minimal parameterization motivated our earlier development of Coupled Hidden Markov Models (CHMMs) to describe interactions between two people, where the interaction parameters are limited to the inner products of the individual Markov chains. The “influence model,” is a generalization of this approach, and describes the connections between many Markov chains as a network of convex combinations of the chains. This allows a simple parameterization in terms of the “influence” each chain has on the others [10, 11, 12].

The influence model has the unique advantage that its steady-state behavior has the same first-order Eigenstructure as the cross-product of all the constituent Markov chains, despite having logarithmically fewer states. As with the Eigenbehavior representation, the influence representation makes it possible to analyze global behavior while avoiding the exponential number of states typical when using other models of interacting individuals or agents.

The graphical model for the influence model is identical to that of the generalized N -chain coupled hidden Markov model (HMM), but there is one very important simplification. Instead of keeping the entire

$P(S_t^i | S_{t-1}^1, \dots, S_{t-1}^N)$, we only keep $P(S_t^i | S_{t-1}^j)$ and approximate the former with:

$$P(S_t^i | S_{t-1}^1, \dots, S_{t-1}^N) = \sum_j \alpha_{ij} P(S_t^i | S_{t-1}^j)$$

In other words, we form our probability for the next state by taking a convex combination of the pair wise conditional probabilities for our next state given our previous state and the neighbors’ previous state. As a result, we only have N $Q \times Q$ tables and N α parameters per chain, resulting in a total of $NQ^2 + N^2$ transition parameters. It is important to realize the ramifications of these factors being constant: intuitively, it means that the model assumes that *how much* we are influenced by a neighbor is constant, but *how* we are influenced by it depends on its state.

In other words, we form our probability for the next state by taking a convex combination of the pairwise conditional probabilities for our next state given our previous state and the neighbors’ previous state. As a result, we only have N $Q \times Q$ tables and N α parameters per chain, resulting in a total of $NQ^2 + N^2$ transition parameters - far fewer parameters than if we tried to model all possible interactions between all of the Markov chains.

This simplification seems reasonable for the domain of human interactions and potentially for many other domains. Furthermore, it gives us a small set of interpretable parameters, the α values, which summarize the interactions between the chains. By estimating these parameters, we can gain an understanding of how much the chains influence each other.

Another benefit of this representation is comes during the analysis the global dynamics of the system. It has been shown in [12] that we can make statements about the recurrent states and the steady state

probabilities of the global system by analyzing the structure of the influence matrix. The advantage of doing such analysis in the domain of human interaction is in understanding how connections between people effect the overall group behavior. How can we manipulate our links to better propagate information or stop the flow of information among the group? If we want consensus among nodes what kind of network graph will help achieve that, i.e. can we identify and modify the recurrent states of the network?

3.4. Influence Modeling

When we use the influence model to analyze the proximity data from our cell phone experiment, we find that clusters of influence in the proximity data map cleanly to our notion of affiliation and friendship. Clustering the daytime influence relationships allowed 96% accuracy at identifying workgroup affiliation, and clustering the evening influence relationships produced 92% accuracy at identifying self-reported ‘close’ friendships [5, 11].

Similar findings were obtained using the badge platform. During a meeting of over 110 Media Lab sponsors, the combination of influence and proximity predicted whether or not two people were affiliated with the same company with 93% accuracy [6].

The ability to accurately characterize peoples’ social relationships by modeling their interpersonal influence (conditional probability structure) means that we can automatically classify the users’ social networks with high accuracy. If we also allow the system to request that the user label relationships that the system thinks are new, are unusual, or where the relationship characterization is uncertain, then we can achieve very high accuracies with very limited user input. Moreover, this

automatic labeling of relationships works even for dynamically changing social networks; its accuracy is largely a function of the amount of observation data available. Thus, for instance, we could accurately label users’ company membership from only a few hours of meeting data.

4. Social Signals

The importance of social displays has been highlighted by the research of Ambady and Rosenthal [13] and its practical ramifications explored in the popular book ‘Blink’ by Malcom Gladwell [14]. In brief, they have shown that people are able to ‘size up’ other people from a very short (e.g., one minute) period of observation, *even when linguistic information is excluded from observation*, and that people use these ‘thin slice’ characterizations of others to quite accurately judge prospects for friendship, work relationship, negotiation, marital prospects, etc. There is something about how we behave that accurately signals the likely future course of our social interactions.

In order to operationalize these findings of social psychologists, we developed automatic, computerized methods for automatically measuring some of the more important types of social signaling [7]. This computer system objectively measures a set of non-linguistic social signal features by looking at ‘tone of voice’ over one minute time periods. These signal features include

- Activity: a measure of the overall amount of speaking/gesturing,
- Emphasis: the amount of modulation of voice and gesture,
- Engagement: a measure derived from timings of speaker turn-taking behavior,
- Mirroring: a measure of mimicry using only short, distinct utterances,

Unlike most speech or gesture research, the goal is to measure and classify speaker interaction rather than trying to puzzle out the speakers' internal state. Consequently we treat vocal prosody and gesture as a longer-term 'motion texture' rather than focusing on individual gestures or accents. Although people are largely unconscious of this type of behavior, other researchers [13, 14] have shown that similar measurements are predictive of infant language development, judgments of empathy, attitude, and even personality development in children.

Using this 'social perception machine' we can "listen in" to the social signals within conversations, while ignoring the words themselves. Following the social psychologists' examples, we reasoned that a test for our ability to automatically measure social signals would be our ability to predict outcomes from 'thin slice' observation of human interactions. Could we predict human behavior without listening to words or knowing about the people involved? We found that after a few minutes of listening in this way, we were able to use linear combinations of these social signal features to accurately *predict* human behaviors such as:

- who would exchange business cards at a meeting;
- which couples would exchange phone numbers at a bar;
- who would come out ahead in a negotiation;
- who was a connector within their work group;

along with a wide range of subjective judgments, including whether or not a person felt a negotiation was honest and fair or a conversation was interesting.

In a typical case using a linear weighting of the social signal features and a three-class linear decision rule (yes, not enough

information, no) the yes/no accuracy is almost 90 percent. Accuracy is typically around 80 percent with a two-class linear decision rule, where we make a decision for every case. More generally, linear predictors based on the measured social signals typically have a correlation of $r=0.65$, ranging from around $r=0.40$ to as much as $r=0.90$. Most experiments involved around 90 participants, typically 25 to 35 years old, with a third being female .

4.1. Signals and information flow

We have been able to distinguish several types of short-term (30 second) display-like behavior patterns that seem to reliably precede important information flow activities such as exchanging personal identifiers [6, 7]. We can name four of the more common displays as 'excitement,' 'freeze,' 'determined,' and 'accommodating,' to pick terms similar to those used in the animal literature, however these 'displays' are really only clusters in the social signal feature data, defined with no direct reference to the semantics these names might suggest.

The 'excitement' display is characterized by a large amount of rapid, highly modulated speech and body movement. The 'freeze' display is characterized by unusually little vocalization and body movement. The 'determined' display is characterized by rapid responses to other vocalizations and little modulation of speech and body movement. The 'friendly' display is characterized by mirroring behavior, above-average listening time, and well-modulated speech and body movement.

In our sponsor meeting data, with more than 110 subjects at each meeting, we observe that the 'excitement' display predicted trading of business contact information with 80% accuracy (equal error

rate). The ‘freeze’ display, when performed in front of a demonstration, predicted requests for additional information with 80% accuracy (the ‘freeze display’ seems to signal mental concentration on the presentation). In a speed dating event, the woman’s display of ‘excitement’ predicted trading of phone numbers with 72% accuracy.

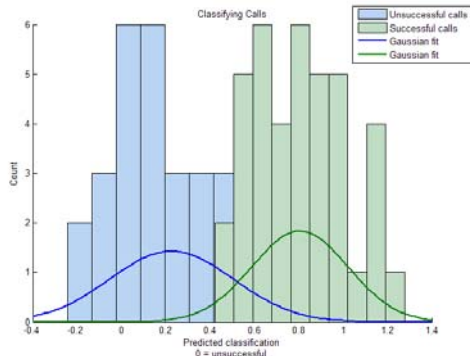


Figure 2. Frequency distribution of successful and unsuccessful calls versus social signaling in a call service center. The cross-validated binary linear decision rule has 87% accuracy.

Figure 2 shows data from a call center, where social signals were measured for 70 calls and compared to the call center ratings of whether or not the calls were successful. The vertical axis is frequency of data points, and the horizontal axis is our prediction based on social signals, with greater values essentially the intensity of the ‘accommodating’ display. A binary decision boundary at 0.45 produces an equal-error accuracy of 87%.

Other experiments have used these social signals to predict ratings of business plan presentations, interest in discussions, outcomes in negotiations, and betweenness centrality in social networks. In each case similar prediction accuracies have been achieved by measurement of social signaling feature statistics. It is especially important to note that these are not classifications of emotion; we know nothing about the subjects’ emotions in these situation. These are classifications of

decision making behavior, as predicted by automatic measurement of the subjects’ social behavior prior to making the decision.

The ability to accurately *predict* peoples’ social interactions by classifying their signaling behavior means that we can automatically classify these interactions with high accuracy. If we also allow the system to request that the user label interactions that the system thinks are new, are unusual, or where the interaction characterization is uncertain, then we can achieve very high accuracies with very limited user input [8].

5. Practical Concerns

Continuous analysis of all interactions within an organization may seem unreasonable, and if misused, could be potentially dangerous. In an attempt to assuage some of these legitimate concerns, several methods of collecting this data will be discussed.

Conversation Postings. In our experiments all the data was stored locally on the individual’s machine. At the end of each day users could potentially review a summary of the number of conversations, the individuals involved, the character of the interaction, and the duration. They could then decide if the data should be shared, private, or permanently deleted. Types of environments where the system might flourish could be places where individuals need to keep careful track of how they spend every minute of their day. Examples of include law firms, and emergency response teams.

Ten Minute Delete / Mute Button: In our experiments each individual’s data device was equipped with an on-off button. A more useful interface might be a button to delete the last ten minutes of the data analysis, or turn off data analysis for ten minutes into

the future. In this way, employees could have a private conversation while at work with a push of the button.

Demanding Environments: In some instances, the environmental demands may supersede privacy concerns. Environments such as these have minimal private conversations, and the needs for all available information is so great, that many of the privacy concerns may not be relevant.

5.1. Related Technologies

Several other projects and technologies have addressed components of the automatic mapping problem. One area that has been widely studied is sensing human context with mobile and wearable devices.

The Active Badge project is one of the earliest examples of a location-aware system that employed infrared (IR) beacons to locate and route phone calls to users [15]. Today's mobile devices have the ability to locate users by cell tower ID, GPS, and scanning fixed beacons such as Bluetooth devices. Projects from both academia and industry are using spatial context to enable location-based services.

By performing repeated scans, mobile devices can measure changing social context of their users. Projects have used both specialized hardware [16, 17, 18] and off-the-shelf smart phones [5, 19, 20] to scan with both IR and radio frequency (RF), e.g. Bluetooth. These social context aware systems are largely intended to support face-to-face collaboration by revealing the user's social context and promoting interaction. Proximity scans have also been used to generate social metadata for images that enables sharing between dyads of proximal people on a per image basis [21].

As described above, our research group has conducted experiments in speed dating, pitching business plans, and conference behavior that have identified behavioral

features useful for the prediction of human interest [6, 22]. Using methods similar to ours, Gatica-Perez et al 2005 [23] found that HMMs built from relatively simple audio and visual features could predict average group interest level ratings for discussion groups.

6. Conclusions

These data make the point that human behavior is much more predictable than is generally thought, and is especially predictable from the behavior of others. This suggests that humans are best thought of as social intelligences rather than independent actors, with individuals best likened to a musician in a jazz quartet. We can predict the behavior of these individuals from that of their associates because they are so attentive and automatically reactive to the surrounding group that they almost cease to be an individual at all.

These data also make it clear that the conditions under which we trade contact information, request information, join groups and so forth, can be quite well *predicted* by location, proximity, and signaling behavior. As a consequence we can analyze behavior using statistical learning tools such as Eigenvector analysis and influence modeling, in order to infer social relationships without needing to understand the detailed linguistic or cognitive structures surrounding social interactions.

We are now using these methods to perform dynamic, continuous analysis of entire commercial organizations [24]. Our initial results suggest that analysis of the face-to-face interactions within the workplace can radically improve the functioning of the organization. By aggregating this information, interpreting it in terms of work tasks, job satisfaction, and productivity, and then modeling the relationship between organizational dynamics and organizational

performance, we are moving towards a new level of performance in the management of complex human organizations.

Acknowledgements

Thanks to all of my collaborators and students for the hard work in forging these tools and this body of data. For future

7. References

1. Allen, T., (1997) Architecture and Communication Among Product Development Engineers, 1997, MIT Press, Cambridge, MA, pp. 1-35
2. Basu, S., Choudhury, T., Clarkson, B., and Pentland, A. (2001) *Towards Measuring Human Interactions in Conversational Settings*. IEEE Int'l Workshop on Cues in Communication (CUES 2001) at CVPR 2001. Kauai, Hawaii.
3. Choudhury, T., and Pentland, A., (2003) *The Sociometer: Measuring Human Networks*, 23d Int'l Social Network Conf., Cancun, MX Feb. 12-16, 2003
4. Pentland, A., (2005) Socially Aware Computation and Communication, IEEE Computer, March, pp. 63-70. See TR 590 <http://hd.media.mit.edu>
5. Eagle, N., Pentland, A. (2003) Social Network Computing, Ubicomp 2003, Springer-Verlag Lecture Notes in Computer Science, No. 2864, pp. 289-296.
6. Gips, J., Pentland, A., "Mapping Human Networks," PerCom '06, Pisa, Italy, March 2006. See TR 586 <http://hd.media.mit.edu>
7. Pentland, A. (2004) Social Dynamics: Signals and Behavior, Int'l Conf. On Developmental Learning, Salk Institute, San Diego, Oct. 20-22. See TR 579 <http://hd.media.mit.edu>
8. Madan, A., Pentland, A., (2006). Vibefones: Socially Aware Mobile Phones, ISWC06, Martigny Switzerland
- collaborators, you will find Matlab code, data, additional information and further publications are available at <http://hd.media.mit.edu> Portions of this paper have appeared in the 5th Int'l Conf. On Development and Learning, Bloomington IL, May 31-June 2, 2006
- Oct. 11-15 See TR 602 <http://hd.media.mit.edu>
9. Eagle, N., and Pentland, A., "Eigenbehaviors: Identifying Structure in Routine, October 2005, See TR 601 <http://hd.media.mit.edu>
10. Pentland, A, Choudhury, T., Eagle, N., Push, S. (2005) Human Dynamics: Computation for Organizations, Pattern Recognition, vol 26, 503-511 See TR 589 <http://hd.media.mit.edu>
11. Dong, W., and Pentland, A., (2006) Multi-sensor data fusion using the influence model, IEEE Body Sensor Networks, April, Boston, MA. See TR 597 <http://hd.media.mit.edu>
12. Asavathiratham, C., The Influence Model: A Tractable Representation for the Dynamics of Networked Markov Chains, in Dept. of EECS. 2000, MIT: Cambridge.
13. Ambady, N., Rosenthal, R. (1992). Thin slices of expressive behavior as predictors of interpersonal consequences: A meta-analysis. *Psychological Bulletin*, 111(2), 256-274.
14. Gladwell, M. (2004) *Blink*. New York: Little Brown
15. Want, R., A. Hopper, et al. (1992). "The active badge location system." *ACM Transactions on Information Systems (TOIS)* 10(1): 91-102.
16. Holmquist, L. E., J. Falk, et al. (1999). "Supporting Group Collaboration with Interpersonal Awareness Devices." *Personal Technologies* 3: 13-21.
17. Iwatani, Y. (1998). "Love: Japanese Style." *Wired News* 11.
18. Terry, M., E. D. Mynatt, et al. (2002). "Social net: using patterns of physical

- proximity over time to infer shared interests.” Conference on Human Factors in Computing Systems: 816-817.
19. Paulos, E. and E. Goodman (2004). “The familiar stranger: anxiety, comfort, and play in public places.” Proceedings of the 2004 conference on Human factors in computing systems: 223-230.
20. Eagle, N. and A. Pentland (2005). “Social serendipity: mobilizing social software.” Pervasive Computing, IEEE 4(2): 28-34. See TR 580 <http://hd.media.mit.edu>
21. Davis, M., S. King, et al. (2004). “From context to content: leveraging context to infer media metadata.” Proceedings of the 12th annual ACM international conference on Multimedia: 188-195.
22. Madan, A., Caneel, R., Pentland, A. (2005) “Voices of Attraction,” AugCog Symposium of HCI 2005, Las Vegas, NV. See TR 584 <http://hd.media.mit.edu>
23. Gatica-Perez, D., I. McCowan, et al. “Detecting Group Interest Level in Meetings.” Proc. IEEE Int. Conf. on Acoustics, Speech and Signal Processing (ICASSP).
24. Ara, K., et al, (2006) “Sensible Organization Inspired by Social Sensor Technologies”, Technical Report No. 602, <http://hd.media.mit.edu>