Improving Public Health and Medicine by use of Reality Mining

A Whitepaper for the Robert Wood Johnson Foundation
Alex Pentland, David Lazer, Devon Brewer, Tracy Heibeck

Executive Summary

We live our lives in digital networks. We wake up in the morning, check our e-mail, make a quick phone call, commute to work, buy lunch. Many of these transactions leave digital breadcrumbs – tiny records of our daily experiences, as illustrated in Figure 1. Reality mining, which pulls together these crumbs using statistical analysis and machine learning methods, offers an increasingly comprehensive picture of our lives, both individually and collectively, with the potential of transforming our understanding of ourselves, our organizations, and our society in a fashion that was barely conceivable just a few years ago. It is for this reason that reality mining was recently identified by *Technology Review* as one of "10 emerging technologies that could change the world" (*Technology Review*, April 2008).

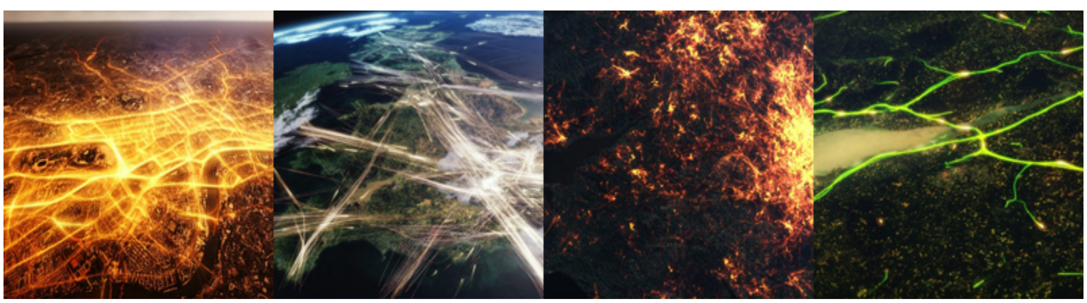

Figure 1: Patterns in human activity: automobiles, airplanes, telephone calls, ships. See http://www.bbc.co.uk/britainfromabove/stories/visualisations/index.shtml for amazing videos of the time evolution of these activity patterns.

Many everyday devices provide the raw database upon which reality mining builds; sensors in mobile phones, cars, security cameras, RFID ('smart card') readers, and others, all allow for the measurement of human physical and social activity. Computational models based on such data could dramatically transform the arenas of both individual and community health. Reality mining can provide new opportunities with respect to diagnosis, patient and treatment monitoring, health services use, surveillance of disease and risk factors, and public health investigation and disease control. The goal of this paper is to survey the potential of reality mining to improve public health and medicine, and to make recommendations for action.

Currently, the single most important source of reality mining data is the ubiquitous mobile phone. Every time a person uses a mobile phone, a few bits of information are left behind. The phone pings the nearest mobile-phone towers, revealing its location. The mobile phone service provider records the duration of the call and the number dialed.

In the near future, mobile phones and other technologies will collect even more information about their users, recording everything from their physical activity to their conversational cadences. While such data pose a potential threat to individual privacy, they also offer great potential value both to individuals and communities. With the aid of data-mining algorithms, these data could shed light on individual patterns of behavior and even on the well-being of communities, creating new ways of improving public health and medicine.

To illustrate, consider two examples of how reality mining may benefit individual health care. By taking advantage of special sensors in mobile phones, such as the microphone or the accelerometers built into newer devices like Apple's iPhone, important diagnostic data can be captured. Clinical pilot data demonstrate that it may be possible to diagnose depression from the way a person talks -- depressed people tend to speak more slowly, a change that speech analysis software on a phone might recognize more readily than friends or family do. Similarly, monitoring a phone's motion sensors can also reveal small changes in gait, which could be an early indicator of ailments such as Parkinson's disease.

Within the next few years reality mining will become more common, thanks in part to the proliferation and increasing sophistication of mobile phones. Many handheld devices now have the processing power of low-end desktop computers, and they can also collect more varied data, due to components such as GPS chips that track location. The Chief Technology Officer of EMC, a large digital storage company, estimates that this sort of personal sensor data will balloon from 10% of all stored information to 90% within the next decade.

While the promise of reality mining is great, the idea of collecting so much personal information naturally raises many questions about privacy. It is crucial that behavior-logging technology not be forced on anyone. But legal statutes are lagging behind data collection capabilities, making it particularly important to begin discussing how the technology will and should be used. Therefore, an additional focus of this paper will be on developing a legal and ethical framework for using reality mining techniques in research, medical care, other health service delivery, and public health surveillance and disease control.

1. CAPABILITIES OF REALITY MINING

To date, the vast majority of research on the human condition has relied on single-shot, self-report data: a yearly census, public polls, focus groups, and the like. Reality mining offers a remarkable, second-by-second picture of both individual and group interactions over extended periods of time, providing dynamic, structural information and rich content. Mobile phone location and movement data, call logs, voice analysis during calls, and email records allow a detailed picture of face-to-face, voice, and digital communication patterns. These patterns can, in turn, afford us a new level of insight into problems of interest to public health and medicine.

1.1 Assessment of Individual Health

The basic functionality of mobile phones consists of the digital signal processing and transmission of the human voice. Advanced mobile phones also have accelerometers, so that they can measure the body movement of their users, and geolocation hardware (both GPS and other methods), so that they can report their users' locations. As a consequence, when users carry around and use their mobile phones they produce a rich characterization of their behavior.

Reality mining of these behavior signals may be correlated to the function of some major brain systems. This statistical behavior analysis therefore provides capabilities that can be thought of as a sort of low-resolution brain scanning technology. Figure 2 illustrates the relationship between brain state and observable behaviors for four types of behavior:

- Arousal of the autonomic nervous system produces changes in activity levels. These changes can be measured by audio or motion sensors, and have been successfully used to screen for depression (Stoltzman, 2006; Sung, Marci, and Pentland 2005; France et al., 2000).

- Tight time-coupling between people's speech or movement (called 'influence') is an indication of attention, since such tight coupling cannot be achieved without attending to and modeling the other person. This 'influence' measure has been successfully used for more than 30 years as a screen for language development problems in pre-verbal infants (Jaffee et al., 2001).

- Unconscious mimicry between people (e.g., reciprocated head nods, posture changes, etc.) is mediated by cortical mirror neurons and is very highly correlated with feelings of empathy and trust. Measurements of mimicry are thus considered to be reliable predictors of trust and empathy (Chartrand and Baugh, 1999), and mimicry has been manipulated to dramatically improve compliance (Bailenson and Yee, 2005).

- Consistency or fluidity of movement or speech production is a well-known measure of cognitive load: novel physical activities or those ‘loaded’ by other mental activity have greater entropy (randomness) than activities that are highly practiced and performed with a singular focus. This relationship has long been used for diagnosis in both psychiatry (Teicher, 1995) and neurology (e.g., Klapper, 2003).

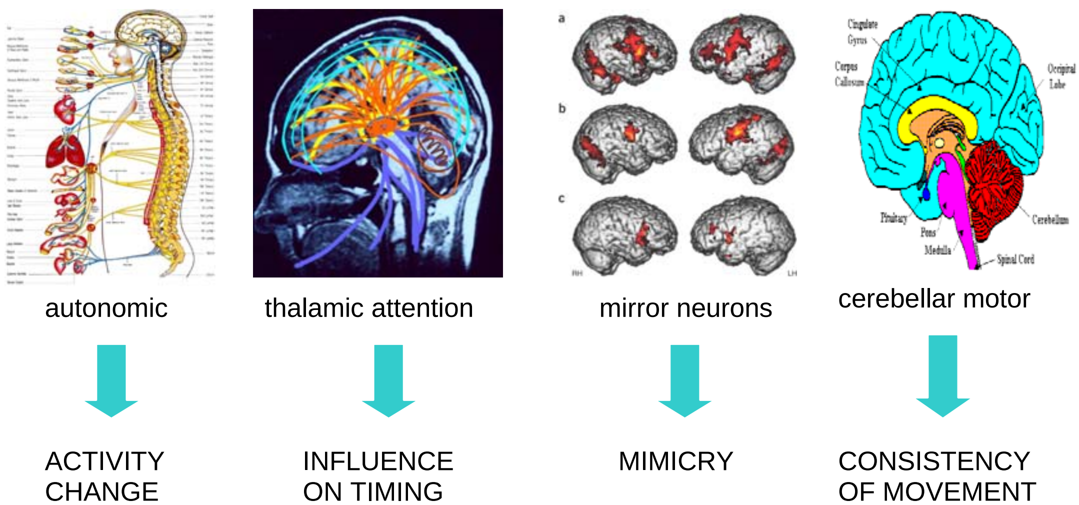

Figure 2: Reality mining has shown that statistical analysis of behavior can be related to the function of some major brain systems, providing capabilities that can be thought of as a sort of low-resolution brain scanning technology.

These qualitative measurements of brain function have been shown to be powerful, predictive measures of human behavior (Pentland, 2008). They play an important role in human social interactions, serving as ‘honest signals’ that provide social cues to dominance, empathy, attention, and trust, and may offer new methods of diagnosis, treatment monitoring, and population health assessments.

Self-report data can also be collected to complement the unobtrusive, automatically-generated and -collected reality mining data streams. The widespread use of portable digital devices such as cell phones and personal digital assistants

(PDAs) enable this marriage of subjective and objective data types. For over a decade, these devices have been used to gather reported data from individuals during the course of their daily lives on such phenomena as symptoms, substance use, and mood.

The technologies for collecting self-reported data in this way are rapidly evolving, but even existing approaches are flexible, automated, and deployable on a large scale. Scheduling of self-reports can be fixed (e.g., on a daily or some more or less frequent timetable), event-based (in response to an event experienced by the respondent or a pattern in one of the data streams objectively recorded by the digital device), or randomly determined. Depending on the device, questions or prompts can be delivered aurally or visually, and responses can be given by voice or touch (e.g., keypad presses or interaction with a display).

Self-reported data offer direct assessment of individuals' cognitive and emotional states, perceptions of events, and information on their behaviors and the contexts in which they are involved that cannot be captured through other reality mining data streams. In many cases, the outcomes of interest in medicine and public health, such as some kinds of symptoms, can be measured only through self-report. By gathering self-reported data in tandem with other reality mining data streams, memory errors are reduced and dynamic aspects of health phenomena are more fully revealed.

1.2 Mapping Social Networks

One of the most important applications of reality mining may be the automatic mapping of social networks (Eagle and Pentland, 2006). In Figure 3 (a), you see a smart phone that is programmed to sense and continuously report on its user's location, who else is nearby, the user's call and SMS patterns, and (with phones that have accelerometers) how the user is moving. One hundred of these phones were deployed to students at MIT during the 2004-2005 academic year. Figure 3 (b) shows the pattern of proximity among the participants during one day; even casual examination shows that the students were part of two separate groups: the Sloan School and the Media Lab.

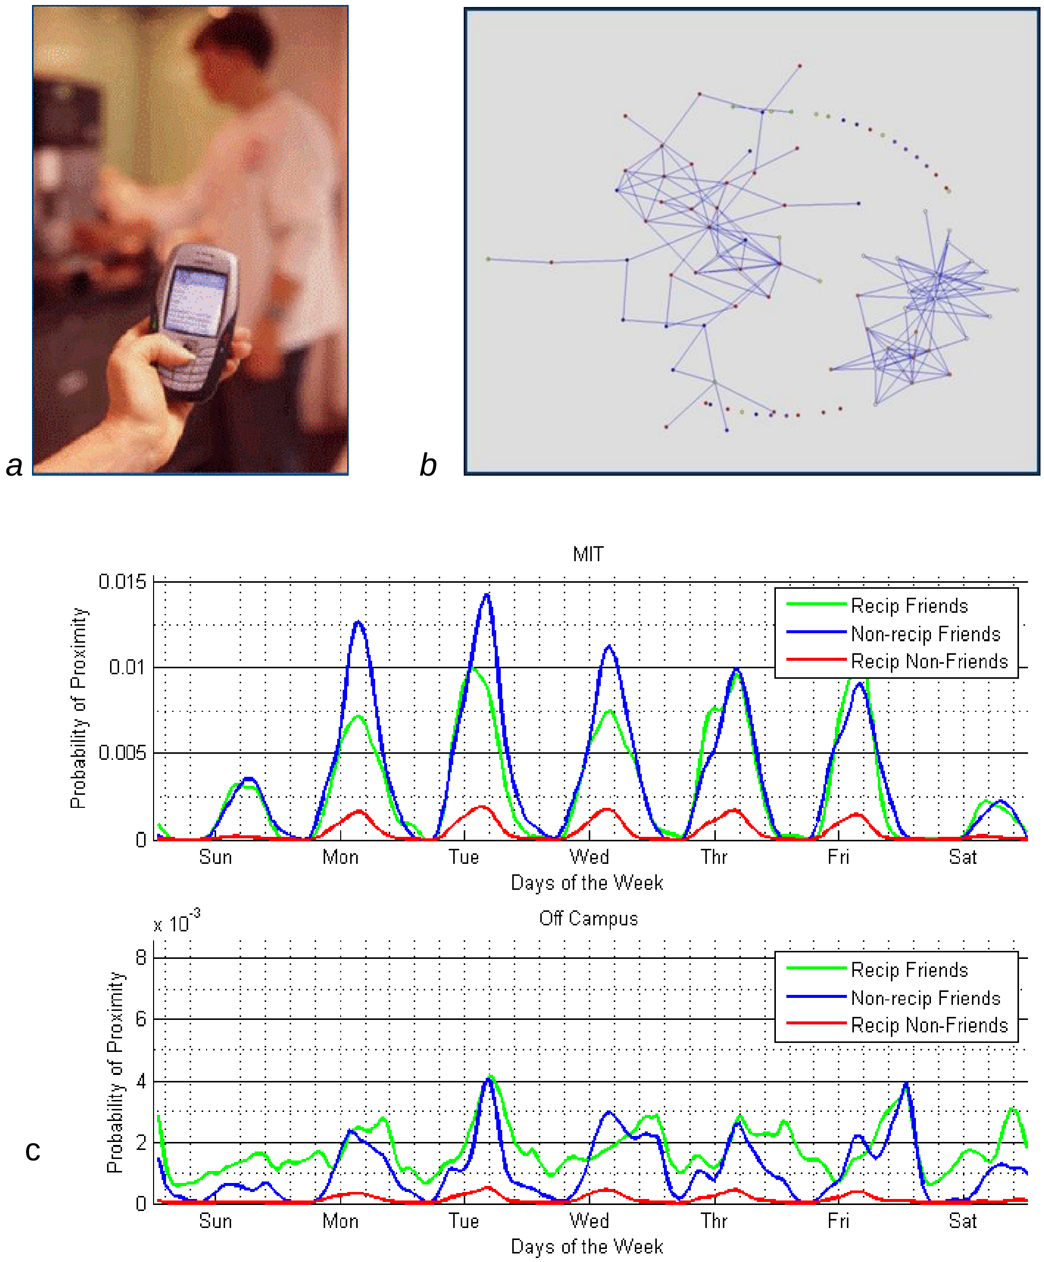

Figure 3: Mapping social networks from mobile phone location / proximity data. 3(a) shows a `smart phone' programmed to sense other people using Bluetooth, 3(b) shows the pattern of proximity between people during one day, and 3(c) shows that different social relationships are associated with different patterns of proximity.

Careful analysis of these data shows different patterns of behavior depending upon the social relationship between people. Figure 3(c) shows the pattern of proximity during one week, and it can be seen that self-reported reciprocal friends (both persons report the other as a friend), non-reciprocal friends (only one of a pair reports the other as a friend), and reciprocal non-friends (neither of a pair reports the other as a friend) exhibit very different patterns (Eagle, Lazer and Pentland, 2007). By using more sophisticated statistical analysis we can map each participant's social network of friends and co-workers with an average accuracy of 96% (Dong and Pentland, 2007).

Reality mining's capability for automatic social network mapping is now being used in a variety of research applications. As an example, a current research project underway at MIT is aimed at understanding health-related behaviors and infectious disease propagation. At this time, we have above 80% participation of students in a MIT dormitory that includes freshmen and upperclassmen, and are beginning to compare the behavior and health changes that freshmen normally experience with the changes in their various social networks. This experiment should help to disentangle causal pathways about how social networks influence obesity and other health-related behaviors, as well as provide unprecedented detail for modeling the spread of infectious disease (see http://hd.media.mit.edu/socially_aware.html).

1.3 Beyond Demographics to Behavior Patterns

Most government health services rely on demographic data to guide service delivery. Demographic characteristics, however, are a relatively poor predictor of individual behavior, and it is behavior – not wealth, age, or place of residence – that is the major determinant of health outcomes. Reality mining provides a way to characterize behavior, and thus provides a classification framework that is more directly relevant to health outcomes (Pentland, 2008).

The pattern of movement between the places a person lives, eats, works, and hangs out are known as a *behavior pattern*. Reality mining research has shown that most people have only a small repertoire of these behavior patterns, and that this small set of behavior patterns accounts for the vast majority of an individual's activity (Pentland, 2007).

The fact that all mobile phones constantly measure their position (either through GPS or by finding the nearest cell tower) means that we can use reality mining of mobile phone location data to directly characterize an individual's set of behavior patterns. We can also cluster together people with similar behavior patterns in order to discover the independent subgroups within a city.

Figure 4(a) shows movement patterns with popular ‘hang outs’ color coded by the different subpopulations that populate these destinations, where the subpopulations are defined by both their demographics and, more importantly, by their *behaviors.* Figure 4(b) shows that the mixing between these different behavior subpopulations is surprisingly small.

Understanding the behavior patterns of different subpopulations and the mixing between them is critical to the delivery of public health services, because different subpopulations have different risk profiles and different attitudes about health-related choices. The use of reality mining to discover these behavior patterns can potentially provide great improvements in health education efforts and behavioral interventions.

a

b

Figure 4: Analysis of travel patterns allows discovery of largely independent subpopulations within a city. Movement patterns (a), measured from GPS mobile phones, allow (b) segmentation of the population into subpopulations with differing behavior patterns, and measurement of the ‘mixing’ between those groups (Sense Networks 2008).

1.4 Assessment of Population Health

Once we have discovered the subpopulations within a community, measured their behavior patterns, and characterized their mixing with other subpopulations, we can then use reality mining techniques to assess their health as a population. Figure 5 shows a plot of regional communication diversity (e.g., the amount of mixing with other subpopulations) and the corresponding index of deprivation. This index is a socioeconomic status measurement that is a combination of metrics such as average income levels, access to healthcare, and education.

This graph shows some of the raw power offered by reality mining: it represents a data set consisting of 250 million hashed (anonymized) phone numbers and 12 billion phone calls. Of particular interest is that the regions with a greater diversity of communication tend to be the least deprived. Diversity of communication behavior has a significant correlation of $r=-.75$, $p<0.001$ with the regional index of deprivation (Eagle, 2008).

These data, drawn from town councils across the UK, show that it is possible to use patterns of communication to identify ‘information ghettos’ that may have serious social problems. This capability may allow government and public health services to be far more responsive to citizen needs than is possible by using census or survey data.

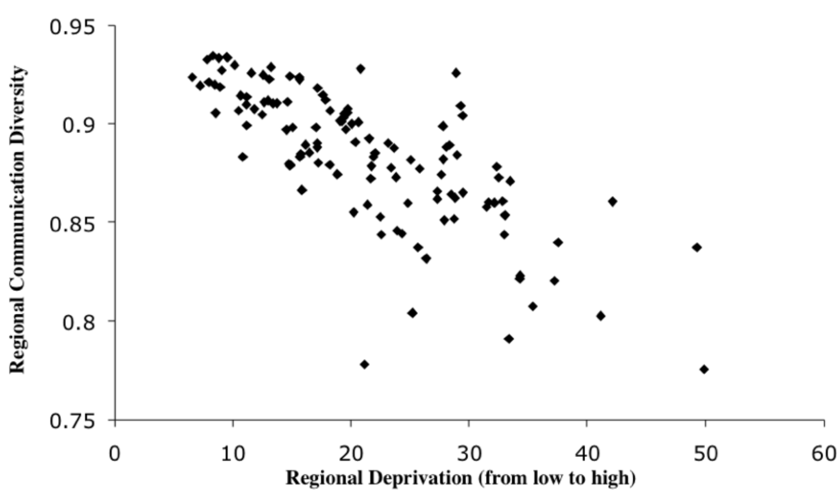

Figure 5. A plot of regional communication diversity and the corresponding index of deprivation.

Other recent technological developments are also emerging to help us better understand population health. For example, Google has developed a method – Google Flu Trends – to detect influenza outbreaks indirectly by tracking the frequency of World Wide Web searches for terms related to influenza-like illnesses (Ginsberg et al., in press). For geographic areas as small as states in the U.S., Google researchers have demonstrated that such search frequencies correlate strongly with estimated influenza incidence based on conventional surveillance of cases detected in a Centers for Disease Control and Prevention (CDC) network of sentinel laboratories and physicians. While the CDC estimates have a reporting lag of 1-2 weeks, the reporting lag for Google Flu Trends estimates is only a day or less. Google Flu Trends may be particularly useful in those areas where good surveillance systems are lacking.

Despite these valuable features, Google Flu Trends does not offer some information provided through traditional surveillance, such as the demographic characteristics of the ill. Google Flu Trends estimates require large populations of Web users, thus estimates for smaller geographic units (e.g., small-to medium-sized cities, towns, and areas) might not be possible. Similarly, Google Trends outbreak detection strategies have not yet succeeded for less common conditions (e.g., enteric infectious diseases). In principle, approaches like Google Flu Trends may be susceptible to bias from search activity unrelated to local disease incidence (such as that stimulated by media coverage of outbreaks in other communities).

Traditional sentinel surveillance efforts can also be improved with current technologies. For instance, the Automated Epidemiologic Geotemporal Integrated Surveillance System (AEGIS), developed by Children's Hospital Boston, involves Internet-based data collection, management, and analysis systems to produce more timely (weekly) estimates of incidence than the CDC effort, and for much smaller geographic areas (sub-county level across Massachusetts). In another example, almost 30,000 residents of three European countries (Belgium, the Netherlands, and Portugal) voluntarily report on their influenza symptoms on a weekly basis at the Gripenet web sites (van Noort et al., 2007). The patterns of estimated incidences of influenza-like illnesses over time from Gripenet correspond closely to those from traditional sentinel surveillance.

2. THE FUTURE POTENTIAL OF REALITY MINING

In the previous section we discussed how reality mining has the potential to assess individual health, to automatically map social networks, to discover subpopulations with different behavior patterns, and to assess the health of those subpopulations. In this section we will explore how these capabilities may facilitate research and public health delivery in areas ranging from encouraging healthy behaviors to monitoring of medical treatments.

2.1 Health Behaviors

Despite compelling evidence, most efforts to encourage healthy behavior and medical compliance continue to be organized around conscious decision making only, neglecting the social dimension almost entirely. By understanding how to leverage social networks, we may achieve more in terms of behavioral change.

For example, research suggests that some chronic health-related conditions/behaviors are "contagious," in the sense that individual-level outcomes are linked to other individuals with whom one shares social connections. Both smoking behavior (Christakis and Fowler, 2007) and obesity (Christakis and Fowler, 2008) have been shown to spread within social networks. Smoking and obesity likely serve as good models for other health related behaviors, such as diet, exercise, general hygiene, and so on.

These findings, however, beg for an examination of the causal mechanism – an essential step if interventions are to be designed to improve public health. For example, is the diffusion of these behaviors and conditions driven by the emergence of norms within the network – e.g., smoking is cool; one should exercise frequently, etc.? Alternatively, is the diffusion driven directly by the social component of the relevant behaviors – e.g., smoking, eating, or exercising with one's friends? The type of data needed to understand the causal mechanism is exactly the fine granularity data that reality mining is poised to provide.

Further, once the causal mechanisms are better understood, reality mining might yield specific points of leverage for effective interventions. For example, if certain behaviors are indeed contagious, this would suggest that targeting individuals in key parts of the network could prove useful (although privacy issues are relevant here; see privacy discussion). Taking this a step further, one could imagine using reality mining to evaluate particular public health interventions. Ideally, program evaluations should test not only whether an intervention was effective, but also the *theory* underlying the intervention. Consider, for example, an intervention based on targeting particular individuals and changing their behaviors. In an attempt to create an avalanche of change, it would be good to know if a given intervention failed because the targeting failed, or because the avalanche failed to materialize despite successful targeting.

2.2 Infectious Disease

With GPS and related technologies, it is increasingly easy to track the movements of people. Mobile phones, in particular, allow the tracing of people's movements and physical proximities over time more precisely than can be done by current approaches (often an individual's recall of their movements) (Gonzalez, Hidalgo, & Barabasi, 2008; Eagle & Pentland, 2006). How might a pathogen, such as the bird flu, driven by physical proximity, spread through a population?

As the world becomes increasingly interconnected through the movement of people and goods, the potential for global pandemics of infectious disease rises as well. In recent years, outbreaks of SARS and other serious infectious diseases in widely separated but socially linked communities highlight the need for fundamental research on disease transmission and effective prevention and control strategies. In developed countries, public health officials typically investigate cases of serious infectious disease (e.g., tuberculosis, SARS, anthrax, measles, Legionnaires' disease, etc.) to identify the source of infections and other cases of disease, and prevent further transmission. Such infections could be transmitted from a point source or person-to-person (directly or indirectly through contaminated airspace, surfaces, or objects).

In any of these scenarios, investigations are difficult and time consuming, while transmission often continues. Logs of location tracking data from cases' cell phones could be examined to identify places where cases might have acquired or transmitted infection, thereby facilitating the investigation. People often forget all the locations they have visited, even for recent periods, and similarly might not know many of the people to whom they were exposed or might have exposed themselves, all of which underlines the potential value of systematically analyzing such records for disease control.

Many ordinary infections already exact a high cost on society. Acute viral respiratory and gastroenteric infections such as the "common cold," influenza, and "stomach flu" produce large negative economic impacts – more than any class of disease – in the United States and elsewhere. It may be possible to determine the relationship between participants' exposures – to symptomatic others and locations visited by symptomatic others – and participants' subsequent respiratory and gastroenteric symptoms. Utilizing reality mining to map social networks may also help to determine whether persons in central social network positions and those with high spatial mobility are more or less vulnerable to infection.

Reality mining tools might also assist in the detection of outbreaks of temporarily disabling disease. For instance, acute illnesses that cause sufferers to reduce their physical activity and mobility (even confining them to bed) or communication behavior, such as influenza, should be noticeable in several types of reality monitoring data streams. At the population level, fluctuations in digital traces of these behaviors may indicate outbreaks of temporarily disabling infectious diseases. To discern outbreaks in this way, long-term data of these sorts would be required to estimate normal background variation. Community-specific unique events (e.g., labor strikes, civil strife, power outages, etc.) would also have to be eliminated as alternate explanations before positing the presence of an outbreak.

2.3 Environmental health

Epidemiologic investigations of the links between individuals' exposures to airborne pollutants, such as particulate matter, carbon monoxide, and nitric oxide, and various health conditions have relied on a variety of exposure measurement methods. To date, most studies of this type have been based on comparisons of aggregates of persons (e.g., residents of a particular neighborhoods or cities, or students at specific schools), with exposure measurements applying to all individuals in a given group. Recently, researchers have begun using more individualized measures of exposures. Some such studies employ static measures of exposure (e.g., estimated pollution levels for a particular residence based on pollution levels recorded from the nearest fixed site sensors in the community and/or inferred from proximity to roadways with known traffic frequencies summarized over time). Measures used in other recent studies provide snapshots of exposure, such as those derived from portable instruments worn by research subjects or biomarkers (e.g., traces of pollutants found in blood, urine, or breath). These latter approaches, while valuable, are labor-intensive and costly, and can be effectively deployed only in small samples for brief observation periods or short series of assessments, thus limiting their widespread use and hampering more detailed understanding of environmental health.

Air pollution levels can vary dramatically over short distances and time scales in urban and other environments, and environmental health experts have called for more precise and dynamic measures of time-activity patterns in relation to exposures. Location tracking data generated by cell phones, when coupled with measurements of ambient air pollution at numerous places in a community (gathered from existing air quality monitoring stations and/or inferred from vehicle traffic patterns and locations of industrial facilities), may offer just the kind of exposure measurement needed. This inexpensive approach would yield dynamic and temporally and spatially more precise measures of exposure suitable for studying large samples of individuals. The location tracking data might even permit differentiating time spent indoors vs. outdoors, through momentary observations in which an individual's (cell phone's) location can be detected from cell tower data but not through GPS data. (GPS readings require a line of sight to satellites overhead and are thus generally not available when indoors.)

2.4 Mental Health

Even though they are quite treatable, mental diseases rank among the top health problems worldwide in terms of cost to society. Major depression, for instance, is the leading cause of disability in established market economies (RAND Corporation, 2004). Reality mining technology might assist in the early detection of psychiatric disorders such as depression, attention deficit hyperactive disorder (ADHD), bipolar disorder, and agoraphobia.

Diagnoses of psychiatric disorders often are based on both subjective states and observable behaviors. In clinical settings, measurement of these states and behaviors for diagnostic purposes is based overwhelmingly on patient self-report, proxy/informant report (e.g., by a teacher or family member), laboratory performance tasks, and/or clinician assessment of patient behavior in clinical settings. Each of these approaches provides useful information, but the agreement between these methods tends to be moderate, which may produce somewhat unreliable diagnoses. Electroencephalograph (EEG) assessments are quite useful in detecting the possible presence of some disorders (e.g., ADHD), although currently they cannot always differentiate diagnoses reliably (i.e., similar EEG patterns can reflect different disorders).

Many signs and symptoms of these kinds of psychiatric disorders explicitly or implicitly relate to an individual's physical movement and activity patterns and communicative behavior, usually with reference to particular temporal periods or cycles. Data streams from reality mining approaches allow direct, continuous, and long term assessment of these patterns and behaviors. Accelerometers in mobile phones, if carried in a pocket, for instance, might reveal fidgeting, pacing, abrupt or frenetic motions, and other small physical movements. Location tracking functions reveal individuals' spatial and geographic ranges, variation in locations visited and routes taken, and overall extent of physical mobility. The frequency and pattern of individuals' communications with others and the content and manner of speech might also reflect key signs of several psychiatric disorders.

The value of these data streams would be multiplied when combined with data reported in real-time by individuals about their psychological states and the contexts in which they are involved. Such reports could be through a variety of response modes on a cell phone and might be triggered by patterns in the location, movement, or communication data or collected on a fixed (hourly, daily, weekly, etc.) or random schedule.

Reality mining methods for diagnosis of medical and psychiatric disorders might be particularly valuable with children and adolescents, who may report their past emotional, mental, and physical states less reliably and articulately than adults. Linguistic and cultural differences between patients, their families, and clinicians also can make diagnosis more difficult, and reality mining approaches might yield diagnostic data that circumvent these challenges to a large degree.

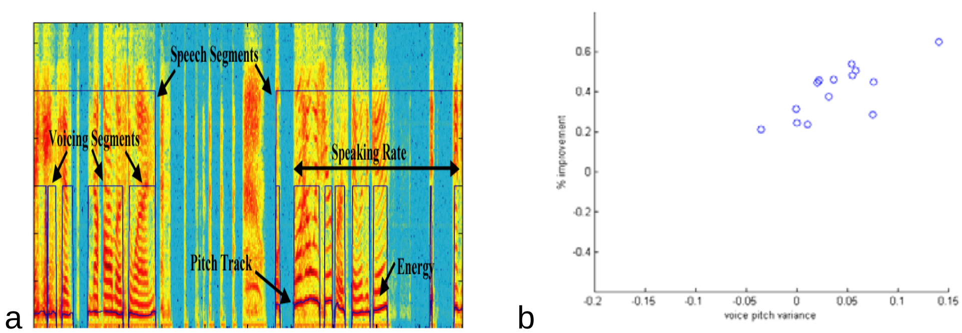

Figure 6: (a) Voice analysis to extract activity, influence, mimicry, and consistency measures. (b) As estimates of depression level, there is a correlation of r = 0.79 between these telephone-based measures and the Hamilton Depression Index.

For a more specific example of the potential power of reality mining technology in aiding diagnosis, consider the data presented in Figure 6. Researchers have long known that speech activity can be affected in pathological states such as depression or mania. Thus, they have used audio features such as fundamental frequency, amplitude modulation, formant structure, and power distribution to distinguish between the speech of normal, depressed, and schizophrenic subjects (France et al., 2000; Stoltzman, 2006). Similarly, movement velocity, range, and frequency have been shown to correlate with depressed mood (Teicher, 1995) (These results can be understood in terms of the qualitative brain system assessment illustrated in Figure 2).

In the past, performing such measurements outside the laboratory was difficult given the required equipment's size and ambient noise. Today, however, even common cell phones have the computational power needed to monitor these correlates of mental state, as illustrated in Figure 6a. We also can use the same methodology for more sophisticated inferences, such as the quantitative characterization of social interactions. The ability to use inexpensive, pervasive computational platforms such as cell phones to monitor these sensitive indicators of psychological state offers the dramatic possibility of early detection of mental problems.

2.5 Treatment Monitoring

Once a course of treatment (whether behavioral, pharmaceutical, or otherwise) has been chosen, it is important for a clinician to monitor the patient's response to treatment. The same types of reality mining data used for diagnosis would also be relevant for monitoring patient response to treatment, especially when such data on the patient are available for a period before diagnosis and can serve as a baseline for comparison. Even when these data streams are not relevant for diagnosis, they might be useful in assessing side effects of treatment, such as reduced mobility, activity, and communicative behavior. Because these data are collected in real-time, a clinician would be able to adjust treatment according to the patient's response, perhaps leading to more effective treatment and preventing more costly office visits.

Continuous monitoring of motor activity, metabolism, and so on can be extremely effective in tailoring medications to the individual. Currently, doctors prescribe medications based on population averages rather than individual characteristics, and they assess patients for the appropriateness of the medication levels only occasionally – and expensively. With such a data-poor system, it is not surprising that medication doses are frequently over- or underestimated and that unforeseen drug interactions occur. Going further, correlating a continuous, rich source of behavioral data to prescription medication use for millions of people could make drug therapies more effective and help medical professionals detect new drug interactions more quickly.

As a more specific example of how reality mining capabilities can be used to monitor treatments, consider the medication needs of Parkinson's patients. To function at their best, Parkinson's patients' medications must be optimally adjusted to the diurnal variation of symptoms. For this to occur, the managing clinician must have an accurate picture of how each patient's combined lack of normal movement (hypokinesia) and disruptive movements (dyskinesia) fluctuates throughout a typical day's activities.

To achieve this, we combined movement data from wearable accelerometers with standard statistical algorithms to classify the movement states of Parkinson's patients and provide a timeline of how those movements fluctuate throughout the day. Two pilot studies were performed, consisting of seven patients, with the goal of assessing the ability to classify hypokinesia, dyskinesia, and bradykinesia (slow movement) based on accelerometer data, clinical observation, and videotaping. Using the patient's diary as the gold standard, the result was highly accurate identification of bradykinesia and hypokinesia. In addition, the studies classified the two most important clinical problems— predicting when the patient "feels off" or is about to experience troublesome dyskinesia—perfectly (Klapper, 2003). This type of fine-grained information, key to monitoring patients' treatment, is a strong endorsement of the value of reality mining techniques.

3. OTHER HEALTH APPLICATIONS THAT LEVERAGE REALITY MINING

The ability of reality mining to automatically map social networks has many potentially important applications in public health. In this section we will explore three key examples:

1) Reinforcing connections to an individual's social support system, since this may be the most effective way to encourage adoption of healthy behavior patterns (Franks et al., 1992).
2) Emphasizing network interactions and leveraging group dynamics as part of a social network approach to health education (Pentland, 2005).
3) Detecting breakdowns in the social support system, so that the support system is more efficient.

3.1 Improving Connection to the Social Support System

The ability to map social networks, and to differentiate between friendship networks, work networks, and family networks, provides the ability to automatically provide better connections to the appropriate social support network. This step is important because *appropriate* reinforcement of social connections may be the most effective way to encourage adoption of healthy behavior patterns (Franks et al., 1992). As an example, consider how elders could be better connected to their social support network through such an approach.

One can imagine a system that strengthens the social ties between an elder patient and their friends by leaving occasional voice mail reminders to call and talk. Such a system could also watch for a marked change in behavior—such as decreased movement, socializing, or voice patterns—and then use its knowledge of the persons' personal networks to reach out and connect them by voice mail reminders. The system need not tell people something is specifically wrong or describe why it left a particular message. Nor would it call the doctor except in extreme circumstances, because doing so could violate elders' privacy and might actually interfere with proper medical support. Instead, this type of monitoring would work to strengthen the social support network when the need is likely to be most significant.

3.2 Leveraging Social Groupings

A second example of leveraging the ability to automatically map social networks is taking a social network approach to health education. Rather than a primary focus on teaching facts, we can instead emphasize network interactions and seek to leverage group dynamics so that *participants take the initiative to actively learn* about health related behaviors by (Pentland, 2005).

Children are particularly well-suited to this social network approach for encouraging self-actuated learning, because they tend to be extremely sensitive to social context. To test this social network concept in reinforcing active learning, we created DiaBetNet, a computer game for young diabetics that leverages smart phone functionality (Kumar and Pentland, 2002). DiaBetNet capitalizes on children's passion for social games to encourage young diabetics to keep track of their food intake, activity, and blood sugar level.

A typical day in the life of a diabetic child using DiaBetNet would unfold as follows. In the morning, the child clips his DiaBetNet case—containing a smart phone and glucose meter—onto his belt and goes off to school. Throughout the day, the smart phone records his activity from the accelerometer, data from the glucose meter about glucose and insulin levels, and user-entered information about food consumption. At any time, the user can see a graph that summarizes the day's activity, carbohydrate consumption, and glucose data. From time to time, a wireless Internet connection sends these data to a secure central server.

DiaBetNet is a group gaming environment that requires guessing blood-sugar levels based on information that wearable sensors collect: the more accurate the answers, the higher the score. For example, imagine that a user named Tom begins to play DiaBetNet with others on the wireless network. Transformed into his cherished alias, Dr. T, Tom finds that his fellow players were all within 30 milligrams per deciliter of guessing their blood sugar levels correctly, but his guess was closer than anyone else's.

Tom challenges a DiaBetNet player called Wizard and looks through Wizard's data. Although Wizard was euglycemic in the morning, he ate a late lunch. Therefore, Tom decides that Wizard's glucose level would be high and guesses 150 mg per dl. Wizard guesses his glucose to be 180 mg per dl. Tom wins again and grabs five more points. He shoots a brief conciliatory message to his vanquished foe and signs off.

In clinical trials, 93 percent of DiaBetNet participants successfully used the system. The Game Group transmitted significantly more glucose values than the Control Group. The Game Group also had significantly less hyperglycemia—glucose 250 mg per dl or greater—than the Control Group. Youth in the Game Group displayed a significant increase in diabetes knowledge over the four-week trial. Finally, more youth in the Game Group monitored their hemoglobin levels.

3.3 Detecting Breakdowns in Supportive Social Networks

As a final example, consider how the automatic detection of breakdowns in the pattern of social support may help make the support system more efficient. Figure 7 presents a schematic drawing of a post-operative ward at a major hospital. We used reality mining methods to monitor this ward for one month, mapping out the networks of communication and interaction among the 70 nurses and doctors.

We then compared these patterns of supportive network interactions to subjective judgments from both nurses and doctors about the availability of required treatment information, and to their judgments about personal stress levels. In both cases we found correlations of approximately r = 0.70 between these subjective factors and deviations from the normal pattern of communication within the supportive network. In light of this tight correlation between subjective perception of difficulty and variations from normal patterns of support network communication, it is not surprising that approximately 80% of the major delays in patient recovery time, and 60% of the failures in patient scheduling, could be detected by monitoring for breakdowns in the patterns of communication and interaction (Olguin, Gloor and Pentland, 2009).

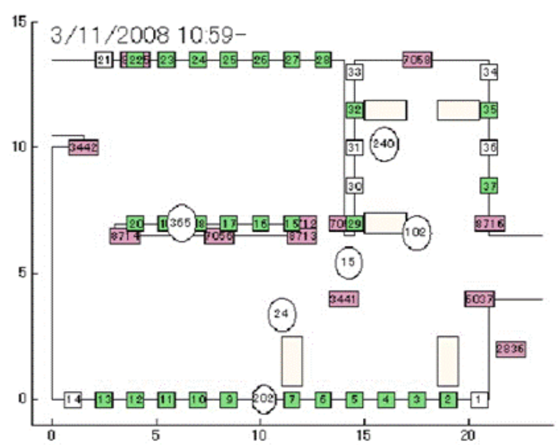

Figure 7: Detecting breakdowns in supportive social networks. This is a real-time schematic of a post-operative ward in a major hospital showing the automatic mapping of behaviors from nurses and doctors using wearable sensors.

4. REALITY MINING AND THE NEW DEAL ON DATA

Reality mining of behavior data is just beginning. In the near future it may be common for smart phones to continuously monitor a person's motor activity, social interactions, sleep patterns, and other health indicators. The system's software can use these data to build a personalized profile of an individual's physical performance and nervous system activation throughout the entire day. If these rich data streams were combined with personal health records, including medical tests taken and the medicines prescribed, there is the possibility of dramatic improvements in health care.

Creating such an information architecture, however, requires safeguards to maintain individual privacy. One approach to this problem is to place control and ownership of as much personal information as possible in the hands of the individual user, a proposal that is central to most proposals for creating personal medical records.

We suggest that a similar approach, a 'new deal' around questions of privacy and data ownership, be taken for data collected using reality mining: individuals own their own data. The simplest approach to defining what it means to 'own your own data' is to go back to Old English Common Law for the three basic tenets of ownership: the rights of possession, use, and disposal.

1. You have a right to *possess* your data. Companies should adopt the role of a Swiss bank account for your data, where you can check your data out whenever you'd like.
2. You, the data owner, must have full control over the *use* of your data. If you're not happy with the way your data is being used, you can remove it.
3. You have a right to *dispose* or *distribute* your data. If you want to destroy it or remove it and redeploy it elsewhere, it's your call.

Social network mapping and the resulting subpopulation information inherently involves other people. As a consequence, some of the thorniest challenges posed by reality mining's ability to sense the pulse of humanity revolve around data access and sharing. There are enormous risks to both individuals and corporations in the sharing of data about individuals. Robust models of collaboration and data sharing, between government, industry and the academy need to be developed; guarding both the privacy of consumers as well as corporations' legitimate competitive interests are vital here. The use of anonymous data should be enforced and analysis at the group level should be preferred over that at the individual level.

Thus, we need to adopt policies that encourage the combination of massive amounts of anonymous data. Aggregate and anonymous location data can produce enormous benefits for society. Patterns of how people move around can be used for early identification of infectious disease outbreaks, protection of the environment, and public safety. It can also help us measure the effectiveness of various government programs, and improve the transparency and accountability of government and non-profit organizations. Thus, advances in analysis of network data must be approached in tandem with understanding how to create value for the producers and owners of the data, while at the same time protecting the public good. Clearly, our notions of privacy and ownership of data need to evolve in order to adapt to these new challenges.

This raises another important question: how do we design institutions to manage the new types of privacy issues that will emerge? It seems likely that new types of institutions are required to deal with this information, but what form should they take? Private companies will also have a key role in this new deal for privacy and ownership. Perhaps market mechanisms can be put in place that allow people to give up their data for monetary or service rewards.

5. RECOMMENDATIONS

In summary, a computational medical and public health science based on reality mining is emerging, a science that leverages the capacity to collect and analyze data with a breadth and depth that was previously inconceivable. The capacity to collect and analyze massive amounts of data has unambiguously transformed such fields as biology and physics. The emergence of such a data driven “computational public health’’ science (CPH science) has been much slower, largely driven by a few intrepid computer scientists, physicists, and social scientists. If one were to look at the leading disciplinary journals in health science and related disciplines, there would be minimal evidence of an emerging CPH science engaged in reality mining of these new kinds of digital traces. What then are the obstacles that stand in the way of a computational public health science?

First of all, there are significant infrastructural barriers to the forward movement of computational health science. The leap from today’s static, snap-shot based public health science to a proactive, dynamic, and CPH science is considerably larger than, for example, from biology to a computational biology, in large part due to the scale and complex ownership of the infrastructure that makes reality mining possible. The resources available to exploratory research in public health are significantly smaller, the computational abilities of public health scientists likely generally lag behind those in the other sciences, and even the physical (and administrative) distance between medicine and public health departments and engineering or computer science departments tends to be greater than for the other sciences.

On the tool side, the availability of easy-to-use tools would greatly magnify the presence of a CPH science. Just as mass-market computer assisted design software revolutionized the engineering world decades ago, common CPH analysis tools and the sharing of data will lead to significant advances. The development of these tools can, in part, piggyback on tools developed in biology, physics, but also requires substantial investments in applications customized to public health needs.

In addition, many important challenges exist on the data side, primarily around access and privacy. Properly managing the issues around privacy is essential. As the recent NRC report on GIS data highlights, it is often possible to pull individual profiles out of even carefully anonymized data. A single dramatic incident involving such a breach of privacy could produce a set of statutes, rules, and prohibitions that could strangle this nascent field in its crib.

What is necessary – now – is to produce a self-regulatory regime of procedures, technologies, and rules that reduces this risk but preserves most of the research potential. As part of this, it is necessary for IRBs to vastly increase their technical knowledge to understand the potential for intrusions and harm to individuals, because the possibilities do not fit their current paradigms for harm. In the longer run, it may be necessary to rethink how IRBs are organized, possibly involving, for example, audits of the safeguards researchers have instituted. These safeguards, in turn, may prove a useful model for industry in enabling their internal research.

To be avoided, however, is either the retreat of an emerging computational public health science into the exclusive domain of private companies, or the development of a “Dead Sea Scrolls” model, with academic researchers sitting on private data from which they produce papers that cannot be critiqued or replicated. Neither scenario will serve the long-term public interest in the agglomeration of knowledge.

Finally, the academic community needs to figure out how to train computational social scientists. A key requirement for CPH analysis to be successful is the development of complementary and synergistic explanations spanning different fields (e.g., epidemiology and network science) and scales (from individuals to small groups and organizations to entire nations). Certainly, in the short run, CPH science needs to be the work of teams of public health and computer scientists. In the longer run, the question will be: should academia be building computational public health scientists, or teams of computationally literate public health scientists and public health literate computer scientists?

The emergence of “cognitive science” in the 1960s and 1970s offers a powerful model for the development of a computational health science. Cognitive science emerged out of the power of the computational metaphor of the human mind. It has involved fields ranging from neurobiology to philosophy to computer science. It attracted the investment of substantial resources to create a common field, and it has created enormous progress for the good in the last generation. We would argue that a computational public health science has a similar potential, and is worthy of similar investments.

References

Bailenson, J., and Yee, N. 2005. Digital chameleons: Automatic assimilation of nonverbal gestures in immersive virtual environments. *Psychological Science, 16(10)*: 814-819.

Chartrand, T., and Bargh, J. 1999. The chameleon effect: The perception-behavior link and social interaction. *J. Personality and Social Psychology*, *76(6)*: 893-910.

Christakis, N., and Fowler, J. 2007. The spread of obesity in a large social network over 32 years. *New England Journal of Medicine*, *357*: 370-379.

Christakis, N., and Fowler, J. 2008. The collective dynamics of smoking in a large social network. *New England Journal of Medicine*, *358*: 2249-2258.

Dong, W., and Pentland, A. 2007. Modeling influence between experts. *Lecture Notes on AI: Special Volume on Human Computing*, *4451*: 170-189.

Eagle, N. Aug. 2008. Behavioral inference across cultures: Using telephones as a cultural lens. *IEEE Intelligent Systems*.

Eagle, N., and Pentland, A. 2006. Reality mining: Sensing complex social systems. *Personal and Ubiquitous Computing*, *10(4)*: 255-268. http://hd.media.mit.edu

Eagle, N., Lazer, D., and Pentland, A. 2007. Inferring friendship from proximity, unpublished manuscript.

France, D., et al. July 2000. Acoustical properties of speech as indicators of depression and suicidal risk. *IEEE Trans. Biomedical Eng.*, 829-837.

Franks, P., Campbell, T.L., and Shields, C.G. Apr. 1992. Social relationships and health: The relative roles of family functioning and social support. *Social Science & Medicine*, 779-788.

Ginsberg, J., Mohebbi, M.H., Patel, R.S., Brammer, L., Smolinski, M.S., and Brilliant, L. (in press). Detecting influenza epidemics using search engine query data. *Nature.*

Gonzalez, M., Hidalgo, C., and Barabasi, A.L. June 5, 2008. Understanding individual human mobility patterns. *Nature*, *453*, 779-782.

Jaffee, J., Beebe, B., Feldstein, S., Crown, C., and Jasnow, M. 2001. Rhythms of dialogue in early infancy. *Monographs of the Society for Research in Child Development, 66(2):* 264.

Klapper, D. 2003. Use of a wearable ambulatory monitor in the classification of movement states in Parkinson's Disease. Master's thesis, Harvard-MIT Health Sciences and Technology Program.

Kumar, V., and Pentland, A. 2002. DiaBetNet: Learning and predicting blood glucose results to optimize glycemic control, 4th Ann. Diabetes Technology Meeting, Atlanta, GA.
www.diabetestechnology.org.

Olguin, D., Gloor, P., and Pentland, A. 2009. Capturing individual and group behavior with wearable sensors, AAAI Spring Symposium, Stanford, CA.

Pentland, A. 2005. Socially aware computation and communication, *IEEE Computer*, *38(3)*: 33-40.

Pentland, A. 2007. Automatic mapping and modeling of human networks. *Physica A: Statistical Mechanics and Its Applications, 378(1)*: 59-67.

Pentland, A. 2008. *Honest signals: How they shape your world.* MIT Press, Cambridge, MA.

Rand Corporation 2004.
http://www.rand.org/pubs/research_briefs/RB9055/index1.html

Sense Networks 2008.
http://www.sensenetworks.com

Stoltzman, W. 2006. Toward a social signaling framework: Activity and emphasis in speech. Master's thesis, MIT EECS.
http://hd.media.mit.edu

Sung, A., Marci, C., and Pentland, A. 2005. Objective physiological and behavioral measures for tracking depression. Technical Report 595, MIT Media Lab.
http://hd.media.mit.edu.

Teicher, M.H. 1995. Actigraphy and motion analysis: New tools for psychiatry. *Harvard Rev. Psychiatry, (3)*: 18-35.

Van Noort, S.P., Muehlen, M., Rebelo de Andrade, H., Koppeschaar, C., Lima Lourenço J.M., and Gomes, M.G. 2007. Gripenet: An internet-based system to monitor influenza-like illness uniformly across Europe. *Eurosurveillance, 12(7)*: pii=722. http://www.eurosurveillance.org/ViewArticle.aspx?ArticleId=722.

BIOGRAPHICAL SKETCHES OF THE AUTHORS

MIT Prof. Alex (Sandy) Pentland is the pioneer of reality mining technology and a leader in ubiquitous information systems. Prof. Pentland is a co-founder of the Center for Future Health at the University of Rochester, and is one of the most-cited computer scientists in the world. He is the founder and Director of Human Dynamics Research within the MIT Media Laboratory.

Prof. David Lazer is Director of the Program on Networked Governance at the Kennedy School of Government, Harvard University, and a leading scholar on networks, who has published in a wide array of leading social scientific and scientific journals.

Dr. Devon Brewer is the director of Interdisciplinary Scientific Research, a research and consulting firm in Seattle, and an affiliate assistant professor at the University of Washington. He conducts research on diverse topics in the social and health sciences, such as research methods and design, statistics, memory, social networks, infectious disease, drug abuse, and crime and violence.

Dr. Tracy Heibeck obtained her PhD in Psychology from Stanford University, and did her clinical training at The Children’s Hospital (Boston). She later became a staff member in Behavioral Medicine at The Children’s Hospital and an Instructor at Harvard Medical School. Dr. Heibeck is also an award-winning technical writer.