

Using Reality Mining to Improve Public Health and Medicine

Alex Pentland, David Lazer, Devon Brewer, Tracy Heibeck
Contact author: pentland@mit.edu, E15-387, The Media Lab, MIT.

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. 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).

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 have the potential to 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 planning, surveillance of disease and risk factors, and public health investigation and disease control.

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 to improve 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 such as 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 – a depressed person tends 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 chapter will be the development of a legal and ethical framework concerning the data used by reality mining techniques.

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.

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 1 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 1: 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.

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 2(a), you see a smart phone that is programmed to sense and report continuously 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 2(b) shows the patterns 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.

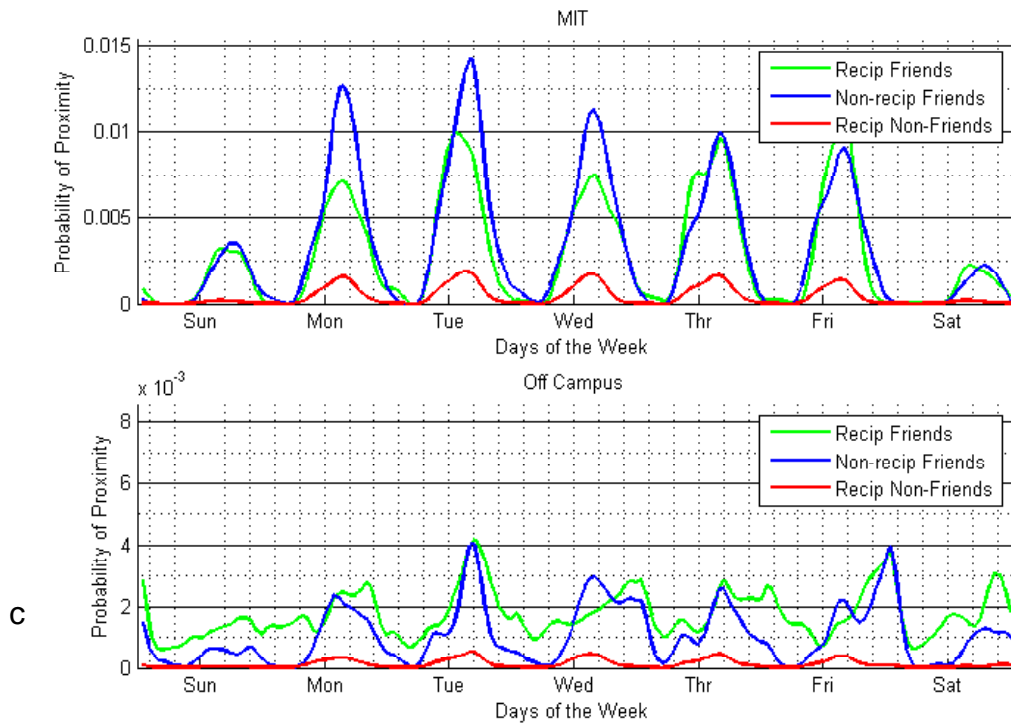



Figure 2: Mapping social networks from mobile phone location/proximity data. 2(a) shows a 'smart phone' programmed to sense other people using Bluetooth, 2(b) shows the pattern of proximity between people during one day, and 2(c) shows that different social relationships are associated with different patterns of proximity.

Careful analysis of these data shows different patterns of behavior depending upon the social relationship between people. Figure 2(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.

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 many 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 population.

Figure 3(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 3(b) shows that the mixing between these different behavior subpopulations is surprisingly small.



3(a)



3(b)

Figure 3: 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).

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.

2. THE FUTURE POTENTIAL OF REALITY MINING

In the previous section, we discussed how reality mining has the potential to assess individual health, to map social networks automatically and to discover subpopulations with different behavior patterns. 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 focused on conscious decision making, 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) seem 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? Or might the apparent spread of these behaviors reflect individuals seeking out others with similar inclinations? The type of data needed to understand the causal mechanism is exactly the fine granularity data that reality mining can 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

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.

With GPS and related technologies, it is increasingly easy to track the movements of people (Gonzalez, Hidalgo, & Barabasi, 2008; Eagle & Pentland, 2006). Logs of location tracking data from cell phones could prove invaluable to public health officials when investigating cases of serious infectious disease (e.g., tuberculosis, SARS, anthrax, measles, Legionnaires' disease, etc.) to help identify the source of infections and prevent further transmission. 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.

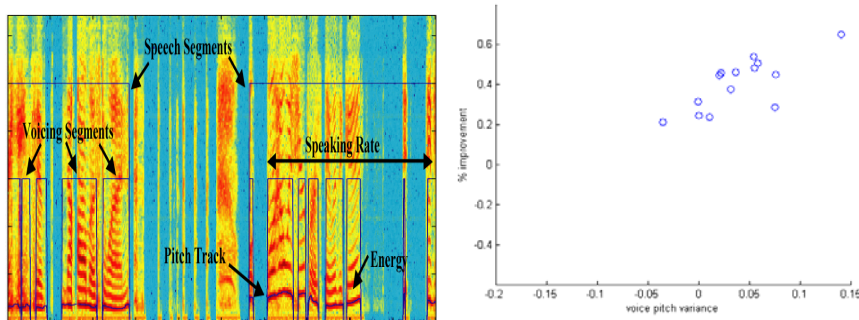
2.3 Mental Health

Even though they are treatable, mental diseases rank among the top health problems worldwide in terms of cost to society. Major depression, for example, 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 hyperactivity disorder (ADHD), bipolar disorder, and agoraphobia.

Many signs and symptoms of these types 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 allow direct, continuous, and long term assessment of these behavior patterns. Accelerometers in mobile phones 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 the 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.

For a more specific example of the potential power of reality mining technology in aiding diagnosis, consider the data presented in Figure 4. 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). Today, common cell phones have the computational power needed to monitor these sensitive indicators of psychological state, offering the possibility of early detection of mental problems.



4(a)

4(b)

Figure 4: (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.

2.4 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. Changes in mobility, activity, and communicative behavior could be collected in real-time, allowing clinicians to adjust treatment according to the patient's response, perhaps leading to more effective treatment and preventing more costly medical visits.

Self-report data can also be collected to complement the unobtrusive, automatically-generated and –collected reality mining data streams. In many cases, the outcomes of interest in medicine and public health (e.g., some kinds of symptoms) can only be measured through self-report. By gathering self-reported data in tandem with other reality mining data streams, memory errors can be reduced and dynamic aspects of health phenomena more fully revealed.

As a more specific example, 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 the day.

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. 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 high accurate identification of bradykinesia and hypokinesia. In addition, the studies classified the two most important 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. 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 self-reports and personal health

records, including medical tests and 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’ for 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, enabling you to 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 are being used, you can remove them.
3. You have a right to *dispose* or *distribute* your data. If you want to destroy them or remove them and redeploy them elsewhere, it’s your decision.

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 of the pulse of humanity concern 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.

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 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. 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.

4. SUMMARY

Reality mining, although still in its infancy, is poised to quickly become more common, due in large part to the rapid proliferation and increasing sophistication of mobile phones. Many mobile phones and other technologies already collect a great deal of information about their users – data such as physical activity and conversational cadences – and this will only increase. Computational models based on such data could dramatically transform many areas of human life. Here, we have focused on improvements that could be realized in individual and community health. Reality mining can provide new opportunities with respect to diagnosis, patient and treatment monitoring, health services planning, surveillance of disease and risk factors, and public health investigation and disease control, and doubtless others, yet unexplored.

In many respects, one of the most important applications of reality mining may be the automatic mapping of social networks. Reality mining's capability for automatic social network mapping is now being used in a variety of research applications and has clear implications for work in infectious disease, health behaviors, mental health, and treatment monitoring. While such data pose a potential threat to individual privacy, they also offer great potential value both to individuals and communities. Current legal statutes are lagging far behind our data collection capabilities, making it particularly important to begin discussing how this technology will and should be used.

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., and Pentland, A. 2006. Reality mining: Sensing complex social systems. *Personal and Ubiquitous Computing*, 10(4): 255-268.
- 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.
- 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.
- 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 measure 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.

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 science and social science 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.

