On The Collective Nature of Human Intelligence

Alex (Sandy) Pentland

Room E15-387, The Media Lab

Massachusetts Institute of Technology

20 Ames St, Cambridge 02139 USA

Pentland@media.mit.edu, 1-617-253-0648

Abstract

A fundamental assumption of cognitive science is that the individual is the correct unit of analysis for understanding human intelligence. I present evidence that this assumption may have limited utility, that the social networks containing the individuals are an important additional unit of analysis, and that this `network intelligence' is significantly mediated by non-linguistic processes. Across a broad range of situations these network effects typically predict 40% or more of the variation in human behavior.

Key Words: theory of mind, cognitive science, evolution, socioscope, network intelligence, individual intelligence.

On The Collective Nature of Human Intelligence

Although we know that some aspects of our thought and culture can be modeled as a function of our social network rather than a function of the individual, a fundamental assumption of cognitive science is that the individual is the fundamental unit of analysis (Goldstone & Janssen 2005, Axelrod, 1977, Putnam 1975, Wegner 1995). We think of ourselves as self-aware individuals with free will, making decisions that shape our lives and carve out our place in the world.

Many influential thinkers go further, supporting the idea that humans begin as a blank slate, and that as a consequence culture floats free of biological constraint (Dengler 1991). To use the metaphor of information processing, they claim that our biology provides a completely general computational platform for running the program of human consciousness, allowing culture to be infinitely malleable. We are individuals, with free will, and rationality, albeit limited.

Researchers, however, have not traditionally accorded the same level of intellectual independence and flexibility to other apes. Although there is important evidence that other apes have cultures (Whiten et al 1999), that these cultures are transmitted at least in part by social learning (Witen 2005), that they possess basic numeracy and (proto)language (Hauser 2005), and at least simple theory of mind (Bryne et al 2004), most human observers still seem to consider ape cognition a thin layer built upon an unconscious animal substrate. Consequently, ape behavior is traditionally thought of in terms of instincts for social displays and responses rather than in terms of conscious cognitive structures and communications. It seems to most humans that the other apes are more `*we'* than `*me*'; they are viewed as social animals whose actions are largely determined by the behavior of others via social and herding instincts.

But is the gap between human apes and the other apes really so wide? As with other apes, the computational substrate provided by our biology places strong constraints on our mental processes and cultural properties. For example, certain mental tasks (such as learning to read) are far more difficult than others (such as learning to speak) (Pinker 2002), and Daniel Brown has compiled a list of almost four hundred traits shared by all known human cultures (Brown 1991) Given our close genetic similarity to other apes, it is possible that human behavior and the behavior of other apes are closely related.

One way to shed light upon the nature of human thought is to apply the tools of biological observation, such as those used to study apes in natural surroundings or in natural experiments such as twin studies. We can imagine an Alien observer with an advanced `socioscope' that can accurately and continuously track the behavior of dozens of humans at a time, recording even the finest scale behaviors with near perfect accuracy. The Alien could then compare human behavior to that of other apes, looking for similarities and differences. Similarities would (by the Occam's Razor argument) be imputed to shared evolution, while differences would be due to uniquely human abilities.

If these observations are to establish the extent to which human thought is independent and general-purpose, we must not *a priori* bring in the generally accepted models of cognitive science to code and interpret observed human behavior (Goldstone & Janssen 2005). If we take as given that human behavior is due to individual conscious thought and a complex, general language, then we risk implicitly building in the assumption that human thought is independent and general-purpose. Instead, we would like to examine the behavioral data in the manner of the field biologist, to see what behaviors can be explained using simple mechanisms that are known to exist in other primates, such as social signaling and herding instincts. Once we have discovered which behaviors cannot be accounted for by these simpler explanations, *then* we can follow cognitive science in making the assumption that the individual is the correct unit of analysis and use the tools of cognitive science to account for the remaining unexplained data.

The importance taking a staged approach to analyzing human behavior can be illustrated by an `alternate universe' thought experiment in which language and conscious thought were not the cause of individual behavior, but instead served only to report the complex interaction of instincts within each individual. In such an alternate universe, the role of language and consciousness would be reduced to rationalizing instinctive behaviors, by telling stories that `explain' decisions made unconsciously.

At least some of our language is this type of rationalization. Experiments with split-brain patients show subjects fluently making such rationalizations for the actions guided by the other half of their brain, and are convinced of the truth and causal nature of their statements even though the separation of the two brain hemispheres makes this impossible (Gazzaniga 2005). In addition, we know that rationalizations are ubiquitous in our language and thinking, particularly when we are aroused, fearful, or angry (Picard 1997, Ariely & Lowenstien 2005).

It seems possible that a significant fraction of our conscious cognition is simply explaining decisions made by unconscious processes; but how large is this fraction? To begin to get at this question, our imagined Alien observers must not begin their investigation by using language to explain behavior. Instead, they must begin with signals and instincts, just as humans do when observing the behavior of other apes.

What might our Alien discover by use of the socioscope? I present four hypotheses:

Hypothesis 1: The majority of human behavior is reactive rather than due to conscious thought processes.

Hypothesis 2: The behavior of other humans is a major and perhaps dominant cue determining unconscious reactive behavior.

Hypothesis 3: The operational equivalent of goals, plans, attitudes, and intentions are reliably and automatically recognizable from non-linguistic behavior.

Hypothesis 4: The network of reactive, unconscious social behavior within a group plays an important role in the groups' fitness.

None of these hypotheses would be terribly controversial if applied to an ape troop. Grooming, for instance, seems instinctual and reactive (hypothesis 1), constitutes an enormous fraction of waking behavior (hypothesis 2), is an important signal of dominance (hypothesis 3), and dominance influences group behavior and fitness (hypothesis 4). Thus the question for our hypothetical Alien is: does the same apply to human behavior? The remainder of this paper will describe human behavior as viewed through a socioscope, and argue that these data provide support for the four hypotheses.

The Socioscope

My students and I have built an approximation of this imaginary socioscope, using mobile telephones, electronic badges, and PDAs (Pentland 2005). My collaborators and I have used this socioscope to track the behavior of graduate students in two divisions of MIT, the business school and the Media Laboratory, a group of 100 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 (Eagle and Pentland 2005).

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 (Gips and Pentland 2006). We have used this platform to obtain data from the more than 110 adults that regularly attend the biannual Media Lab sponsor meetings, 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, ½ from North America, and 1/6 from Europe.

The third part of the socioscope consists of a microphone, optional body-worn camera to record the wearers' context, 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 (Pentland 2004, Choudhury & Pentland 2004, Clarkson 2003). This part of the socioscope can be used with audio data from the smart phone, audio from the badge, or (more commonly) audio from body-worn microphones during semi-structured interactions such as speed dating, focus group interviews, or negotiations.

Together, these three sensor platforms have allowed us to observe gross behavior (location, proximity) continuously over months, to more accurately observe behavior (location, proximity, body motion) over one-day periods, and to analyze vocalization statistics with an accuracy of tenths of seconds. The behavioral data are then subject to four main types of analysis: characterization of individual and group distribution and variability, conditional probability relationships between individual behaviors (which I will call `influence'), accuracy of prediction (with equal type I and II error rates), and finally the relationship of these behavioral measures to standard cognitive and cultural metrics.

The Data

Variability in behavior

Our imaginary Alien would see one main daily pattern: subjects leaving their sleeping place to congregate in one building for 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 then back their sleeping place. Variations from this pattern can be broken into principal components, with the top three principal components typically accounting for 80% of the variance across subjects (Eagle and Pentland 2006). Individual subjects typically have a characteristic mix of these three components, accounting for approximately 90% of the variance in their behavior.

In human terms, these three components could be thought of as the weekend pattern, the working late pattern, and the socializing pattern. Even though we are considering largely unmarried young people who are still in school, it seems that there is limited variability in behavior.

If the Alien increased the resolution of the socioscope, so that body motion within a location was also measured, the behavior would be broken into a string of `situated behaviors', such as walking down a hall or sitting in a room (Clarkson 2003). By clustering these data by similarity of both motion and surrounding, the subjects' daily behavior would be broken into strings of 30 common situated behaviors, repeated with variations from day to day. Transitions from one behavior to the next could be predicted with a typical perplexity (branching factor) of four, although with many rare choices being possible. That is, the Alien could build a Markov model of the daily behavior with 30 states and an average branching factor of four between states, and using this model correctly predict the subjects future behavior with greater than 50% accuracy. If similar `situated behaviors' are aggregated into only 10 states, then the prediction accuracy increases to 75%. In human terms, these 30 `situated behaviors' can be cleanly mapped to things like `sitting in a meeting,' `walking along a street,' `eating in a restaurant,' and `shopping in a store.' The small branching factor between successive states shows that the storyline of our lives is limited to a relatively small number of variations.

These data are consistent with the first hypothesis. The data show that our subjects' behavior has quite limited variability, and fairly predictable from the surrounding static, physical context, as would be expected if most behavior was reactive.

Network Influence

The previous data illustrate the stereotypical patterns and limited variability that our imaginary Alien would observe in *individual* subjects. Next, let us ask what behavioral structure the Alien would observe *between* subjects.

Conditional probability relationships between subjects, which I will refer to as influence, allow us to predict the behavior of a subject from the other subjects' data. 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 (Dong & Pentland 2006). Overall, influence (conditional probability relations) between subjects allowed 95% of the variance in personal proximity data to be accounted for by the surrounding network of proximity data.

In human terms, 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 `close friendships' between subjects (as determined by self-report in an on-line survey). On a finer scale, when we looked at influence and proximity during our biannual meetings, we found 93% accuracy at predicting whether or not two people were affiliated with the same company (Gips & Pentland 2006).

On our finest scale, one may measure influence in vocalization. For instance, if I stop talking, how likely are you to immediately start talking, and vice versa. Using 1700 hours of vocalization data from 21 subjects, we found that a person's average influence on the vocalization of others had a correlation of r=0.90 with the social network property of betweeness centrality (a measure of how much a person connects otherwise isolated groups) (Choudhury & Pentland 2004). In human terms, the more someone exhibits the *network* property of being a `connector,' the more they drive the pattern turn-taking in conversations. Moreover, the people in this study were unaware of either their network status or their turn-taking influence.

These data speak to hypothesis two. Our data show that human behavior is surprisingly predictable from the behavior of others in their social network, and much more predictable than when considering only the static physical surroundings. The very large percentage of variance accounted for adds to the plausibility of hypothesis one.

Social Displays

The importance of social displays has been highlighted by the research of Ambady and Rosenthal (1992) and its practical ramifications explored in the popular book `Blink' by Malcom Gladwell (2004). 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.

Looking at the finest scale data, our hypothetical Alien would be able to distinguish several types of `social display', defined as short-term (30 second) display-like behavior patterns that reliably precede important functional activities such as exchanging personal identifiers. For instance, he might notice that in some circumstances the frequency and energy of vocal and body gesturing is a good predictor of subsequent behavior. Our Alien might name four of the more common displays as `excitement', `freeze', `determination', and `accommodation', to pick terms similar to those used the animal literature (Whiten et al 1999, Whiten 2005, Hauser 2005, De Waal 2005). 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 ‘accommodating’ display is characterized by mirroring behavior, above-average listening time, and well-modulated speech and body movement. Remember, however, that these `displays’ are really only distinguished clusters in behavior data, defined with no direct reference to the semantics these names might suggest. The displays are described in more detail in Pentland (2004).

In our sponsor meeting data, with more than 110 subjects at each meeting, our Alien would observe that the `excitement’ display predicted trading of business cards with 80% accuracy (equal error rate), thus the presence of signaling accounted for about 40% of the variance in business card trading behavior (Gips & Pentland 2006). The `freeze’ display, when performed in front of a demonstration, predicted requests for additional information with 80% accuracy, so again signaling accounted for about 40% of the variance. In human terms, the `freeze display’ signals mental concentration on the demonstration.

In a speed dating event, the woman’s display of `excitement’ predicted *both* parties’ trading of phone numbers with 72% accuracy. Here signaling accounted for almost 40% of the variance in dating and contact trading behavior (Madan & Pentland 2006).

In a salary negotiation (conducted for grade in a business school negotiation class), the proportion of `determination' displays by the higher-status participant predicted 36% of variation in salary (Curhan and Pentland 2007). The proportion of `accommodation' displays by the lower status participant predicted 25% of variation in salary. If the displays of both were considered, 40% of the variation in salary could be predicted.

In addition to dating, salary negotiations, networking, and centrality in your social network, this type of social signaling also predicts objective outcomes in sales pitches, bluffing in poker, job hiring, treatment for depression, and subjective judgments of trust, interest, determination, and similar relationship characteristics (Pentland 2004, Sung & Pentland 2005, Stolzman 2006). These data speak to hypothesis three: operational equivalents of goals, plans, attitudes, and intentions are reliably recognizable from simple social signals. Unconscious reactions to these signals provides a possible mechanism for achieving the social reactivity of hypotheses one and two.

What Does It Mean?

If these data were collected from ape troops, and we altered the semantically-loaded labels a bit (e.g., `forage' instead of `work', `food access' instead of `salary'), they would feel entirely unsurprising. We might even argue that the data support the view that the observed behavior is largely automatic, determined by instincts for herding and social display / response pairings, since typically 40% or more of the variation in behavior is determined by the non-linguistic behavior of the surrounding individuals. But, of course, these data are from humanity's best and brightest, and include important behaviors like getting a date, a job, or a raise….so how are we to interpret them?

The first point is simply that human behavior is much more predictable than is generally thought…and remember that MIT graduate research laboratories are famously unregimented and informal. The behavior of most people is likely to be far more regular and predictable. It is also clear that our behavior, including important acts such as dating, hiring, negotiation, group membership and so forth, can be quite well *predicted* by our location, proximity, and signaling behavior. As a consequence, much of our behavior can be *explained* by the non-linguistic behavior of our associates, by invoking simple instincts for herding and social display/response and without explanatory recourse to linguistic or conscious cognitive structures.

There is still an important question about the causality of these social cues in determining behavior. Could the causal chain be through language rather than through unconscious reaction to social context? The presence of similar reactive social structures in other apes means that the simplest, most conservative interpretation is that causality is due to unconscious reaction rather than conscious, linguistic interaction.

However, the other possibility is even more intriguing: let us assume that causality is through the linguistic channel, and the signals we have measured are simply emergent and epiphenomenal. That would mean that our language has simple, accurate, and surface-level cues to goals, intentions, and attitudes. After all, we have measured these surface-level signals and found them quite accurate in their prediction of subsequent behavior. It is hard to believe that people would fail to unconsciously pick up the association between such simple signals and such important outcomes and then put the correlations to good use. That, of course, is pretty much the same thing as saying that much of human behavior consists of unconscious reactions to social signals.

The data support the view that our conscious cognitive structures serve less of a causal role in determining our behavior than is commonly believed in the cognitive sciences (as per Goldstone & Janssen 2005), and may often serve only to record and communicate decisions made by our unconscious and instinctual brain. This does not, of course, argue against the idea that our cognitive structures can override our instincts, only against the idea that almost all human behavior is caused by conscious, cognitive processes.

Elevating the importance of our instinctive social nature relative to the roles of language and conscious cognition feels dangerous to some people, as if we were admitting that we are ruled by some base animal nature. The data do not support such a negative view. Instead, they emphasize that our behavior is deeply and immediately connected with that of other humans. Humans are seen to be social animals, where individuals are best likened to a musician in a jazz quartet. Of *course* we can predict the behavior of these individuals from that of their associates: they are so focused on the group's overall performance and so sensitive to exactly complementing the others in the group that they almost cease to be an individual at all. From this perspective, the data support the view that this immersion of self in the surrounding social network is the *typical* human condition, rather than being isolated examples found in exceptional circumstances.

Using a computational metaphor, one might say that humans have a specialized processor with many built-in functions and hard-wired defaults. On top of that is a powerful programming language that coerces the processor to do its bidding. Some things are hard to do, of course, because of limitations and peculiarities of the processor.

On the other hand, the program can also make use of the specialized hardware functions to make some functions easier. One of those specialized functions is our theory of mind and our ability to effortlessly understand higher order intentionality (roughly, the ability to predict what others will think). Another specialized function is our ability to quickly and unconsciously read and display social signals that communicate internal states such as interest, determination, pleasure, and friendliness. It is these latter sorts of special `hardware' abilities that allow us to coordinate smoothly with the surrounding network of individuals. They are also the abilities that allow a jazz quartet, a Ranger unit, and many other types of human groupings to function better as a collective than as a set of isolated individuals.

Wisdom of the Network

The evidence presented so far supports the plausibility of hypotheses one, two, and three: that the behavior of humans in groups forms a reactive network substantially mediated by unconscious social signals. We now turn to hypothesis four, the claim that this reactive social network plays a role in the fitness of the group. An example will prove illustrative and at the same time advance the argument. Imagine a tribe on the African veldt:

> *Each day the adults go out gathering and hunting, and in the evening return to recount the events and observations of the day, and discuss what to do tomorrow. During the group*

discussion social signaling…tone of voice, laughter, body posture, gesture, and so forth...reflects each individual's desires and interests as well as their position in the social hierarchy. This signaling accompanies each discussion item and the collective social signaling communicates back to each member of the tribe what the group thinks: Is this item new information? Is it important? Does a proposed decision violate vested interests? Does it open up new opportunities for some? At the end, collective decisions have been made, often without explicit declaration, but because each individual knows the sense of the group, they know what duties they must perform during the next day.

Decision making behavior of exactly this kind is commonly observed in ape troupes (Hauser 2005, De Waal 2005), where it is obviously non-linguistic. The same sorts of social signaling are a major mechanism in human group decision making (Brown 1986, Chen et al 2004, Wu et al 2004).

These data address our fourth hypothesis, but perhaps the most interesting property of socially mediated group decision making is that it can be smarter than any of the group's individual members. This power stems from the group's potential to integrate information gathered by many different members. For instance, Chen et al (2003) have shown that a Bayesian aggregation of information guided by social information can reliably outperform all the individual experts contributing to the aggregate prediction. Accurate aggregation of information requires that you pay attention to which opinions are independent of the others so that you don't double count. To do that, you have to keep track of how information flows within the group. You also have to weight each independent opinion by the track record of success for the person giving the opinion, and by their risk profile.

Similar mechanisms have been surveyed by James Surowiecki in his book `Wisdom of Crowds' (2004). Similarly, the *failure* of group decision-making due to social phenomena like groupthink, polarization, etc., is at its heart a failure to keep account of which bits of information are independent and which are copies propagated through the social network, so that information is aggregated inappropriately (Brown 1986, Chen et al 2004).

One may ask if humans can use our talents for tracking social interactions to accurately aggregate information across a group, and thereby produce judgments that are more accurate than those of isolated individual humans? Or, to phrase it more provocatively, can we humans use our special-purpose `social brain' to produce a super-human collective intelligence?

To test this hypothesis, we constructed an experiment in which six groups of three subjects each conducted a treasure hunt, looking for specially-marked images of familiar people (Chilongo 2006). We then asked the subjects how many people had appeared in the background of the target images…something they had not been instructed to attend to. This sort of exercise is analogous to a hunting party that is asked about unusual conditions after returning from a hunt, or a management team that is asked to estimate threats and opportunities that are outside of normal operating procedures. This sort of serendipitous information gathering and assessment is often central to strategic planning.

What the data from this experiment show is that when we aggregate information by taking into account objective social data ---- social network structure as measured by the socioscope during the treasure hunt and risk profiles observed in a separate exercise --- we get estimates of the true distribution that are twice as good as any individual, and four times better than the average individual. Moreover, when we aggregate information using the subjects' *subjective* estimates of social information, we get essentially the same estimates as when using *objective* social information…so apparently people were able to keep track of information flow during the treasure hunt.

This sort of social network awareness is a huge computational task, yet one that humans do effortlessly, perhaps because our computational engine is specialized for just this sort of political thinking. Indeed, there is strong evidence that the computational task of understanding interactions within social networks accounts almost entirely for the rapid expansion of primate frontal cortex (Dunbar 2004).

If you can keep track of all the social interactions and figure out which bits of information are independent, then your group can be reliably smarter than any one individual (Chen et al 2003, Chilongo 2006). One can imagine that groups that are able to accurately integrate information will feel that they have access to a wisdom that transcends human intelligence…and they will be correct. Such an ability would confer an evolutionary advantage, one that depends upon and accrues to the *entire* cultural group regardless of genetic heritage, and it has been argued that this type of evolutionary advantage may be the root cause for the development of religious practice (Wilson 2002).

Conclusion

The evidence presented strongly supports the position that humans must be understood as social animals as well as individuals, and that our behavior and thought processes should be understood as due to non-linguistic network interactions as well as individual properties and dynamics. These data therefore have the potential to broadly alter thinking in the cognitive sciences, since they undermine the assumption that the individual is the correct unit of analysis. These data could also unsettle parts of the social sciences, which have tended to treat culture as isolated from the properties of the individual.

Some might counter that all this is obvious, for we have always known that people learn from other people and use others as convenient repositories of knowledge (Wegner 1995). But I am making a stronger point: that important parts of our personal cognitive processes are *caused* by the network via unconscious and automatic processes such as signaling and imitation, and that consequently, important parts of our intelligence depend upon *network properties*. Human intelligence is in both the individual and the social network, and across a broad range of tasks between 40% and 80% of the variation in human behavior is attributable to non-linguistic interactions with the surrounding network.

This change in perspective could have important practical ramifications. Consider the management of companies, government agencies, and other adult organizations. To improve information aggregation and decision making groupings would be heterogeneous (to spread the reach of available information), grow slowly from a small initial team (making it easier to learn the propagation of opinion within the social network), and be limited in size to less than 150 individuals (the approximate limit on our ability to keep track of social relations) (Dunbar 2004). Similar changes might apply to the organization of childhood learning. Rather than focusing on teaching cognitive skills, one might instead emphasize network interactions (Harris 1995).

Another consequence of network intelligence and unconscious influence on personal cognition might be network effects for explicitly cognitive tasks. For instance, some studies have found that many of our opinions are surprisingly predictable from our associates' opinions, even when the opinion flies in the face of strong counterfactual evidence (Milgram 2004, Bikhchandani et al 1998, Lazer 2001, Chartrand and Bargh 1999). Likewise, the similarity of self-descriptions on personal web pages and similarity of consumer preferences have both been found to fall off exponentially with increasing social network distance (Wu et al 2004).

Such homogeneity of thought seems to be due to continual and largely unconscious enculturation to your immediate social network. One need only think of Milgram's experiments (2004) or Zimbardos' Stanford Prison Experiment (Haney et al 1973) to appreciate the speed and power of unconscious enculturation. It is a widely held suspicion that the disappointing performance of teleconferencing systems and computer software for cooperative work is due to their failure to adequately convey the social signals that mediate this enculturation (Cassell 2000).

By properly channeling and leveraging our human `network intelligence,' we can improve information aggregation and decision making, as illustrated by our `treasure hunt' experiment. There is the potential to dramatically improve the practice of science, the management of organizations, and political governance.

Experimental data, computer code, and papers with additional detail can be found at http://hd.media.mit.edu

References

Ambady, N. & Rosenthal, R. (1992). Thin slices of expressive behavior as predictors of interpersonal consequences: A meta-analysis. *Psychological Bulletin, 111(2), 256-274.*

Ariely, D., & Lowenstien, G., (2005) In the Heat of the Moment: The Effect Sexual Arousal on Sexual Decision Making, *J. Behavioral Decision Making, 18: 1-12*

Axelrod, R., (1977) The dissemination of culture: a model with local convergence and global polarization, *J. Conflict Resolution, 42, (1) 56-76,*

Bikhchandani, S., Hirshleifer, D., & Welch, I., (1998) Learning from the Behavior of Others: Conformity, Fads, and Informational Cascades, *J. Economic Perspectives, 12: 151-170,*

De Waal, F. (2005) *Our Inner Ape*, New York: Riverhead

Brown, R. (1986) *Group Polarization, in Social Psychology (2d Edition)*, New York: Free Press

Brown, D. E. (1991) *Human Universals*, New York: McGraw-Hill

Byrne, R., Barnard, P., Davidson, I., Janik, V., McGrew, W., Miklosi, A., & Wiessner, P., (2004) Understanding Culture Across Species, *Trends in Cognitive Science, 8 (8) 341-346*

Cassell, J. (2000). More than Just Another Pretty Face: Embodied Conversational Interface Agents. *Communications of the ACM 43(4): 70-78* .

Chen, K-Y, Fine, L., and Huberman, B., (2004) Predicting the Future, *Information Systems Frontiers,* 5 (1) 47-61

Chen, K-Y, Fine, L, & Huberman, B, (2003) Eliminating Public Knowledge Biases in Information-Aggregation Mechanisms, *Mangement Science, 50 (7) 983-994*

Chilongo, D. M. (2006) *Social Network Effects on Information Aggregation,* MEng Thesis, MIT EECS Department. Advisor: A. Pentland. See http://hd.media.mit.edu

Clarkson, B., (2003) *Life Patterns,* MIT Phd Thesis, Media Arts and Sciences Program, Advisor: A Pentland. See http://hd.media.mit.edu

Chartrand, T., & Bargh, J., (1999) The Chameleon Effect: The Perception-Behavior Link and Social Interaction, *J. Personality and Social Psychology, 76 (6) 893-910*

Choudhury, T., & Pentland, A., (2004) Characterizing Social Networks using the Sociometer. *Proceedings of the North American Association of Computational Social and Organizational Science.* June 10-12, Pittsburg, Pennsylvania. See http://hd.media.mit.edu

Curhan, J., & Pentland, A., (2007) Thin Slices of Negotiation: Predicting outcomes from conversational dynamics within the first five minutes, to appear *J, Applied Psychology* See http://hd.media.mit.edu

Dengler, C., (1991) *In Search of Human Nature: The decline and revival of Darwinism in American social thought*, New York: Oxford University Press

Dong, W., & Pentland, A., (2006) Multi-sensor Data Fusion Using the Influence Model. *IEEE Body Sensor Networks, pp. 72-75*, April 3-5, Boston, MA. See http://hd.media.mit.edu

Dunbar, R. (2004) *The Human Story*. London: Faber & Faber.

Eagle, N., & Pentland, A. (2006) Reality Mining: Sensing Complex Social Systems, *Personal and Ubiquitous Computing, 10 (4) 255-268.* See http://hd.media.mit.edu

Eagle, N., &Pentland, A., (2006) Eigenbehaviors: Identifying Structure in Routine, *MIT Media Lab Technical Report 596,* See http://hd.media.mit.edu

Gazzaniga, M, (2005) *The Ethical Brain*, Univ. Chicago Press

Gips, J., & Pentland, A., (2006) Mapping Human Networks, *IEEE International Conference on Pervasive Computing and Communications*, Pisa, Italy, March 13-17, 2006. See http://hd.media.mit.edu

Gladwell, M. (2004) *Blink.* New York: Little Brown

Goldstone, R., & Janssen, M., (2005) Computational Models of Collective Behavior, *Trends in Cognitive Science,* 9(9): 424-430

Haney, C., Banks, W. & Zimbardo, P. (1973) Intrapersonal Dynamics in a Simulated Prison, *Int'l J. Criminology and Penology, 1: 69-97*

Harris, J., (1995) Where is the child's environment? A group socialization theory of development, *Psychological Review, 102: 458-489*

Hauser, M., (2005) "Our Chimpanzee Mind," *Nature, 437: 60-63*

Jaffe, J., Beebe, B., Feldstein, S., Crown, C. L., & Jasnow, M. (2001). Rhythms of dialogue in early infancy. *Monographs of the Society for Research in Child Development, 66(2), No. 264.*

Lazer, David, (2001) The Co-Evolution of Individual and Network, *Journal of Mathematical Sociology, 25(1):69-108*

Madan, A., & Pentland, A., (2006) VibeFones: Socially Aware Mobile Phones. *10th International Symposium on Wearable Computing*, Oct. 11-14, Montreux, Switzerland, see http://hd.media.mit.edu

Milgram, S. (2004) *Obediance to Authority: An Experimental View,* New York: HarperCollins

Pentland, A. (2004) Social Dynamics: Signals and Behavior, *Int'l Conf. On Developmental Learning,* Salk Institute, San Diego, Oct. 20-22. See http://hd.media.mit.edu

Pentland, A., (2005) Socially Aware Computation and Communication, *IEEE Computer, 38 (3) 33-40*

Picard, R. (1997), *Affective Computing,* Cambridge: MIT Press,

Pinker, S.,(2002) *The Blank Slate: the Modern Denial of Human Nature,* New York: Penguin Books

Putnam, H., (1975) *The Meaning of Meaning,* Philo. Papers, Vol. 2: Mind, Language, and Reality. Cambridge Univ. Press,

Stolzman, W., (2006) *Toward a Social Signaling Framework: Activity and Emphasis in Speech*, MIT EECS MEng Thesis, Advisor: A. Pentland. See http://hd.media.mit.edu

Sung, M., and Pentland, A., (2005) *PokerMetrics: Stress and Lie Detection,* MIT Human Dynamics Technical Report 594. See http://hd.media.mit.edu

Surowiecki, J, (2004) *The Wisdom of Crowds*, New York: Doubleday

Wegner, D., (1995) A Computer Network Model of Human Transactive Memory, *Social Cognition*, 13(3) 319-339

Wilson, D., (2002) Darwins' Cathedral: Evolution, Religion, and The Nature of Society, Chicago: University of Chicago Press

Witen, A., Goodall, J., McGrew, W., Nishida, T., Reynolds, V., Sugiyama, Y., Tutin, C., Wrangham, R., & Boesch, C. (1999) Cultures in Chimpanzees *Nature, 399: 682-685*

Witen, A., (2005) The Second Inheritance System of Chimpanzees and Humans, *Nature, 437: 52-55*

Wu, F., Huberman, B., Adamic, L, & Tylers, J. (2004) Information Flow in Social Group, *Physica A, Vol. 337, pp. 327-335*

Author Note

Alex (Sandy) Pentland, Media Laboratory, Massachusetts Institute of Technology

The author gratefully acknowledges former graduate students Tanzeem Choudhury, Sumit Basu, Ron Caneel, David Chilongo, Brian Clarkson, Wen Dong, Nathan Eagle, Jon Gips, Anmol Madan, Will Stolzman, and Mike Sung for all their hard work and dedication in building the socioscope and collecting the data, collaborators Joost Bonsen (who named the `socioscope'), Jared Curhan, David Lazer, Carl Marci, and M. C. Martin for many good discussions, and Tracy Heibeck, Deb Roy, and Chris Csikszentmihályi for insightful comments on earlier drafts of this paper. An early version of this paper appeared in the IEEE Int'l Conference on Development and Learning, and portions of the methods and data sections have appeared in IEEE Computer .

Biography

Prof. Alex (Sandy) Pentland is a pioneer in wearable computers, health systems, smart environments, and technology for developing countries. He is one of the most-cited computer scientists in the world.

He is a co-founder of the Wearable Computing research community, the Autonomous Mental Development research community, the Center for Future Health, and was the founding director of the Media Lab Asia. He was formerly the Academic Head of the MIT Media Laboratory, and is MIT's Toshiba Professor of Media Arts and Sciences, and Director of Human Dynamics Research.

He has won numerous international awards in the Arts, Sciences and Engineering.

Correspondence concerning this article should be addressed to Alex Pentland, The Media Laboratory, MIT, Room E15-397, 20 Ames St., Cambridge, MA 02139, USA. Email: Pentland@mit.edu