

Are We One? On the Nature of Human Intelligence

Alex (Sandy) Pentland

Room E15-387, The Media Lab, MIT

20 Ames St., Cambridge, 02139

pentland@mit.edu

Abstract – A fundamental assumption of cognitive science is that the individual is the correct unit of analysis. 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 the distribution of intelligence across individuals in the social network is significantly mediated by non-linguistic processes. Evidence is presented that about 40% of the variance in human behavior may be attributable to non-linguistic social context.

Index Terms – individual, theory of mind, cognitive science, evolution, socioscope.

I. INTRODUCTION

A fundamental assumption of cognitive science is that the individual is the fundamental unit of analysis 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 [1,2,3]. 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 [4]. 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 “one,” separate individuals with free will.

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 [5], that these cultures are transmitted at least in part by social learning [6], that they possess basic numeracy and (proto)language [7], and at least simple theory of mind [8], 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 instead of in terms of cognitive structures and communications. It seems to most humans that the other apes are more ‘we’ than ‘one,’ 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? Daniel Brown, for instance, has compiled a

list of almost four hundred traits shared by all known human cultures [9]. Similarly, certain mental tasks (such as learning to read) are more difficult than others (such as learning to speak) [10]. We also know that our cognitive processes change dramatically when we are angry, afraid, or aroused [11, 12]. These facts support the view that the computational substrate provided by our biology is not completely general, and that its uneven computational abilities biases our mental processes and hence cultural properties.

One way to shed light upon the questions of independence and generality 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 advanced ‘socioscope’ [13] 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.

If these observations are to address the individuality and generality of human thought, however, we must not *a priori* assume that humans are conscious individuals with a complex and general language. If we bring in the generally accepted model of language to code and interpret observed human behavior, then we risk implicitly building in the assumption that human thought is independent and general. 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 of this staged approach to analyzing human behavior can be illustrated by an ‘alternate universe’ thought experiment in which we might find that all human behavior could be explained by some simple function of other humans non-linguistic behavior. In such an alternate universe the role of language and consciousness would be reduced to rationalizing instinctive behaviors, allowing individuals to tell stories to each other that explain the complex interaction of the instincts within them. Language and perhaps consciousness would cease to be causal.

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 [14]. In addition we know that rationalizations in general are ubiquitous in our language and thinking, particularly when we are aroused, fearful, or angry [11, 12].

Let us therefore add to properties of our imagined Alien observers the limitation that they don't have the idea of linear, sequential language, so that it won't occur to them to consider language as an explanation for human behavior. What might such an observer find? They might find that human behavior is more automatic and predictable than is generally thought. They might also find that our actions are better predicted by our social context than by unobserved internal states and linguistic variables. But let us first examine the data.

II. THE SOCIOSCOPE

My students and I have built an approximation of this imaginary socioscope, using mobile telephones, electronic badges, and PDAs [15-20]. 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 350,000 hours of data covering the behavior of 81 people for a period of nine months.

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. This badge platform provides more fine-grained data than the smart phone platform, but has the constraint that it only works within the Media Lab and the batteries only last for one day. 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, 1/2 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. 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 allow 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.

III. THE DATA

A. *Variability in behavior*

Our imaginary Alien would see one main daily pattern, that of 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. Individual subjects typically have a characteristic mix of these three components, accounting for up to 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 unattached young people who are still in school, it seems that there is limited variability in our 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. 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.

B. 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. Overall, influence 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 90% accuracy at identifying self-reported 'close friendships.'

* The extent to which these two networks overlapped predicted reported job satisfaction with a correlation of $r=0.87$.

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.

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 persons average influence on the vocalization of others had a correlation of $r=0.90$ with the social network property of betweenness centrality (a measure of how much a person connects otherwise isolated groups). 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 unconscious of either their network status or their turn-taking influence.

C. Social Displays

The importance of such social displays has been highlighted by the research of Ambady and Rosenthal [21] and its practical ramifications explored in the popular book 'Blink' by Malcom Gladwell [22]. 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. Our Alien might name four of the more common displays as 'excitement', 'freeze', 'dominance', and 'non-aggression', to pick terms similar to those used the animal literature. 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.

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. 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 presentation.

In a speed dating event, the woman's display of 'excitement' predicted trading of phone numbers with 72% accuracy. The display of 'non-aggression' by both parties predicted the trading of business contact information with 78% accuracy. Here signaling accounted for almost 40% of the variance in dating and contact trading behavior.

In a salary negotiation (conducted for grade in a business school negotiation class), the proportion of 'dominance' displays by the higher-status participant predicted 27% of variation in salary. The proportion of 'non-aggression' displays by the lower status participant predicted 30% of variation in salary. If the displays of both were considered, 40% of the variation in salary could be predicted.

IV. 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 supports the view that the observed behavior is largely automatic, determined by instincts for herding and social display / response pairings, since 40% 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, and 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 cognitive structures.

This result supports the view that our cognitive structures serve less of a causal role in determining our behavior than is commonly believed, 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.

The somewhat dismal feel associated with of this view of humanity, however, stems more from the implicit assumption of the individual human as the causal element than from elevating the importance of our basic instinctual nature. We can alternatively take the view of the human as social animal, where individuals are best likened to a musician in a jazz quartet or a Army Ranger living for months with just his eight-man Ranger unit. 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 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.

V. WISDOM OF THE NETWORK

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? 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.

One of the most interesting properties of such a group is that it can be smarter than any of its' individual members. This power stems from the group's potential to integrate information gathered by many different members. In a recent paper, Chen et al [23, 24] showed that aggregating information using a simple modification of Bayes' theorem could reliably outperform all the individual experts contributing to the aggregate prediction. Similar mechanisms have been developed by Malone for governance of organizations, and surveyed by James Surowiecki in his book 'Wisdom of Crowds' [25].

The basic idea is that you weight each independent opinion by the track record of the person giving the opinion, and their risk profile. Because keeping a sense of each individuals' success and caution is important, there may be a temptation for individuals to use social signals to emphasize these properties. But the real catch is that you have to pay attention to which opinions are independent of the others so that you don't double count. 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 [24, 26, 27].

To figure out which information is independent requires, by my count, being able to understand fifth-order intentionality across all the relationships within the group...and perhaps fortunately this is exactly the human limit for intentionality understanding. That is, you need to be able to figure out what a third party will think of an interaction between two others (e.g., I want everyone to agree to this proposal, so I will ask her to heap praise on any man who voices support, so that it is more likely that the other men will also agree).

This is a huge computational task, yet one that humans do effortlessly, apparently 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 [28].

The prize is that 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. 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. One can speculate that such ability would confer an evolutionary advantage[†] and might be related to the development of religious practice.

VI. CONCLUSION

I believe that the evidence I have 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 alter thinking in the cognitive sciences, since they have been built on 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 [3]. But I am making a stronger point: that important parts of our intelligence reside in *network properties*, not individual properties, and that important parts of our personal cognitive processes are *caused* by the network via unconscious and automatic processes such as signaling and imitation (herding). As a consequence, human intelligence is in both the individual and the social network, with as much as 40% of the 'intellectual labor' attributable to non-linguistic signals from 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 higher-order intentionality relations), and be limited in size to less than 150 individuals (the approximate limit on our ability to keep track of social relations) [28]. Similar changes might apply to the organization of childhood learning. Rather

[†] perhaps even a *group selection* advantage, if one believes the arguments of 'top-down causality,' since the survival advantage accrues to the entire cultural group and not just closely related individuals.

than focusing on teaching cognitive skills, one might instead emphasize network interactions [29].

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 [30, 31, 32]. 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 [27].

Such homogeneity of thought seems to be due to continuous and largely unconscious enculturation to your immediate social network. One need only think of Milgrams' experiments [33], Zimbardo's Stanford Prison Experiment [34] 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 [35].

By properly channeling and leveraging our human 'network intelligence' we can improve information aggregation and decision making. 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>

ACKNOWLEDGMENT

Thanks to Joost Bensen for many good discussions, to Tracy Heibeck, Deb Roy, Chris Csikszentmihályi for insightful comments on earlier drafts of this paper .

REFERENCES

- [1] Goldstone, R., Janssen, M., "Computational Models of Collective Behavior," *Trends in Cognitive Science*, Vol 375, July 2005.
- [2] Axelrod, R., "The dissemination of culture: a model with local convergence and global polarization," *J. Conflict Resolution*, Vol. 42, No. 1, pp. 56-76, 1977
- [3] Putnam, H., "The Meaning of Meaning," *Philo. Papers*, Vol. 2: Mind, Language, and Reality. Cambridge Univ. Press, 1975
- [4] Dengler, C., "In Search of Human Nature: The decline and revival of Darwinism in American social thought," New York: Oxford University Press
- [5] Whiten, A., Goodall, J., McGrew, W., Nishida, T., Reynolds, V., Sugiyama, Y., Tutin, C., Wrangham, R., Boesch, C. "Cultures in Chimpanzees," *Nature*, Vol. 399, pp. 682-685, 1999.
- [6] Witen, A., "The Second Inheritance System of Chimpanzees and Humans," *Nature*, Vol. 437, pp. 52-55, September 2005
- [7] Hauser, M., "Our Chimpanzee Mind," *Nature*, Vol. 437, pp. 60-63, September 2005
- [8] Byrne, R., Barnard, P., Davidson, I., Janik, V., McGrew, W., Miklosi, A., Wiessner, P., "Understanding Culture Across Species," *Trends in Cognitive Science*, Vol. 8, No. 8, pp. 341-346, August 2004
- [9] Brown, *Human Universals*, Cambridge: MIT Press, 1991
- [10] Pinker, S., "The Blank Slate: the Modern Denial of Human Nature," New York: Penguin Books, 2002
- [11] Picard, R. (1997), *Affective Computing*, MIT Press, Cambridge, 1997,
- [12] Ariely, D., and Lowenstein, G., "In the Heat of the Moment: The Effect Sexual Arousal on Sexual Decision Making," *J. Behavioral Decision Making*, 18: 1-12 (2005)
- [13] Bensen, J. *Personal communication*
- [14] Gazzaniga, M., "The Ethical Brain," Univ. Chicago Press, 2005
- [15] Clarkson, B., "Life Patterns," MIT Phd Thesis, Media Arts and Sciences Program, October 2003. Advisor: A Pentland. See <http://hd.media.mit.edu>
- [16] Choudhury, T., "Sensing and Modeling Human Networks." Phd thesis, Dept of MAS, MIT, 2003. Advisor: A. Pentland. See <http://hd.media.mit.edu>
- [17] Eagle, N., Pentland, A. (2003) *Social Network Computing*, Ubicomp 2003, Springer-Verlag Lecture Notes in Computer Science, No. 2864, pp. 289-296. See <http://hd.media.mit.edu>
- [18] 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>
- [19] Eagle, N., and Pentland, A., "Eigenbehaviors: Identifying Structure in Routine," MIT Media Lab Technical Report 596, October 2005, See <http://hd.media.mit.edu>
- [20] Gips, J., Pentland, A., "Mapping Human Networks," to appear *PerCom '06*, Pisa, Italy, March 2006. See <http://hd.media.mit.edu>
- [21] Ambady, N., Rosenthal, R. (1992). Thin slices of expressive behavior as predictors of interpersonal consequences: A meta-analysis. *Psychological Bulletin*, 111(2), 256-274.
- [22] Gladwell, M. (2004) *Blink*. New York: Little Brown
- [23] Chen, K-Y, Fine, L., and Huberman, B., "Predicting the Future," *Information Systems Frontiers*, Vol. 5, No. 1, pp. 47-61, 2004
- [24] Chen, K-Y, Fine, L, Huberman, B, "Eliminating Public Knowledge Biases in Information-Aggregation Mechanisms," *Management Science*, Vol. 50, No. 7, pp. 983-994, July 2004
- [25] Surowiecki, J, "The Wisdom of Crowds," New York: Doubleday, 2004
- [26] Brown, R. (1986) *Group Polarization*, in *Social Psychology* (2d Edition), New York: Free Press
- [27] Wu, F., Huberman, B., Adamic, L, Tylers, J. "Information Flow in Social Group," *Physica A*, Vol. 337, pp. 327-335, 2004
- [28] Dunbar, R., "The Human Story." London: Faber & Faber. 2004
- [29] Harris, J., "Where is the child's environment? A group socialization theory of development," *Psychological Review*, Vol. 102 pp. 458-489
- [30] Bikhchandani, S., Hirshleifer, D., and Welch, I., "Learning from the Behavior of Others: Conformity, Fads, and Informational Cascades," *J. Economic Perspectives*, Vol. 12, pp. 151-170, 1998
- [31] Lazer, David, "The Co-Evolution of Individual and Network." *Journal of Mathematical Sociology*. 2001; 25(1):69-108
- [32] Chartrand, T., and Bargh, J., (1999) The Chameleon Effect: The Perception-Behavior Link and Social Interaction, *J. Personality and Social Psychology*, Vol. 76, No. 6, 893-910
- [33] Milgram, S. "Obedience to Authority: An Experimental View," New York: HarperCollins, 2004
- [34] Haney, C., Banks, W. and Zimbardo, P. "Intrapersonal Dynamics in a Simulated Prison," *Int'l J. Criminology and Penology*, Vol. 1, pp. 69-97, 1973
- [35] Cassell, J. (2000). "More than Just Another Pretty Face: Embodied Conversational Interface Agents." *Communications of the ACM* 43(4): 70-78 .

