HCI Workshop
Reality Testing: HCI Challenges in Non-Traditional Environments
Montreal, Canada
April 22, 2006

Names & Affiliation:

Anthony P. Glascock, Ph.D.
Professor of Anthropology
Department of Culture & Communication
Drexel University
Philadelphia, PA 19104
USA

David Kutzik, Ph.D.
Associate Professor of Sociology
Department of Culture & Communication
Drexel University
Philadelphia, PA 19104
USA

Short Bio Sketches:

Dr. Glascock is currently Professor of Anthropology in the Department of Culture and Communication at Drexel University. His current research interests focus on the use of technology to improve the quality of life of older adults by allowing them to remain in their own residences. This interest grew out of extensive research on home health care conducted over the last decade in Ireland, Canada and the United States. He has published over 60 articles and book chapters and his latest book, The Aging Experience: Diversity and Commonality Across Cultures was awarded the Richard Kalish Innovative Publication Award presented by the Behavioral and Social Science Section of the Gerontological Society of America. He is President of Behavioral Informatics, Inc., a member of the Board of Directors of Living Independently, Inc. and holds three patents (two US, one Canadian) and has 11 patents pending on a behavioral monitoring system.

Dr. David M. Kutzik is an Associate Professor of Sociology in the Department of Culture and Communication at Drexel University, Philadelphia, PA. Specializing in gerontology, his research interests focus on developing and integrating appropriate technologies into care models designed to maximize independence and functional health of elderly and chronically ill populations. Originally engaged in applied research on elderly health behaviors and long term care delivery case management, Dr. Kutzik played a leading role in the development of non obtrusive behavioral monitoring systems which have lead to a variety of patent protected technologies as well as products. He has authored and coauthored dozens of articles, chapters and presentations on topics ranging from medication adherence and cognitive impairment to assistive technologies for the elderly. Dr. Kutzik is also Vice President of Behavioral Informatics, Inc., Media, PA and a member of the Medical Advisory Board of Living Independently Group, Inc., New York, New York.

# Embrace the Chaos, It's Not Noise: Lessons Learned from Non-Traditional Environments

Although it is possible to test whether a piece of equipment works in a laboratory, it's not possible to test whether a technology works within a confined laboratory environment. This is because a technology is not equipment alone, but instead a technology is the equipment plus its application: that is, what it does in the real world. However, the problem with the real world is that it is messy; it is inhabited with people who have real needs and real problems and who want a technology that meets these needs and solves their problems in a way that makes sense to them and at a cost in time and money that they can afford. Therefore, it is not surprising that developers of new technology are often reluctant to take the technology into this real world where they lose control and a sense of security and where the chaos of real life rears its head and confounds test protocols, schedules and projected costs. "Why can't people behave the way they are supposed to and use the technology the way it was intended", is a phrase often heard when technology is taken from the lab to the real world. Thus, developers spend inordinate amounts of time trying to eliminate the noise of real life, rather than accepting that the noise is the key to making the technology work.

We have spent the last five years dealing with this noise as we moved a new behavioral monitoring technology from the laboratory to the real world. Initially, we viewed this noise as a hindrance to this successful transition, but over time we became convinced that without embracing the chaos and listening very carefully to the noise, our monitoring system could not be successfully moved from the laboratory to the real world. This process has been painful at times, but we believe that the lessons we have learned can benefit other developers and perhaps, help them embrace the chaos also.

Over six years of laboratory research and development on behavioral monitoring technology resulted in the first real world ready iteration of the Everyday Living Monitoring System (ELMS). The ELMS was comprised of five motion sensors and a base station connected, via the Internet, to a website that processes the sensor data and converts them to information that is then displayed with graphics and text for caregivers. The PIN secure website provides a daily summary for each person being monitored for six activities: waking time; bathroom falls; the taking of medication; meal preparation; overall level of activity; and nighttime bathroom use, as well as ambient temperature. In addition the ELMS provides emergency alerts for bathroom falls, non-wake-up and high or low temperature and has the ability to produce monthly summary charts for all monitored activities.

We developed a five stage real world testing design that would allow us to install the system into increasingly more complex environments. In this way, it was hoped, that we would be able, at each stage, to refine the system before moving to the next more challenging environment. We have progressed through four of these stages and our hopes and fears have been more than met as the messiness of the real world has constantly introduced noise and chaos into our well thought out testing protocols. The four environments into which we have tested the ELMS are: 1) an activity of daily living suite at a large regional hospital; 2) the homes of volunteers as part of individual sequential testing; 3) the residences of five elderly individuals in a high rise apartment building with limited care provision; and 4) in the independent residences of 34 chronically ill individuals with nursing services provided by a home health care company. We are just beginning the fifth stage of testing that will compare the effectiveness of the ELMS in six different care environments ranging from assisted living to independent households and include over 200 individuals (The Caring Home Initiative). Although we have learned many things from each of the completed stages, we will limit ourselves to the discussion of the two most illustrative examples from each of the stages.

Stage 1: ELMS was installed in an activity of daily living suite comprised of a one bedroom apartment with a full kitchen, bedroom, living/dining area and bathroom at a large regional hospital. Elderly patients spent the night in the suite in order to be assessed as to their ability to function independently. ELMS, along with a video camera, was installed in the suite to record activity during the overnight stays The two most illustrative findings were: first, do not place sensors near heating and air conditioning vents; and second, less is often more. Because it appeared to be the best location to record movement between the bedroom and bathroom, a sensor was initially placed within 12 inches of a heating and air conditioner vent. The result was real noise as approximately every 15-20 minutes there was a flurry of activity recorded which appeared to indicate that the individual being monitored entered and left the bathroom repeatedly. Review of the video tape indicated no such activity and eventually it was determined that it was the air coming from the vent that triggered the sensor.

Initially, every appliance—microwave, dishwasher, stove, television—and every cabinet and drawer in the kitchen was monitored. Also, multiple motion detectors were placed throughout the living/dining room, bedroom, bathroom and kitchen. The result of this was that much of the data proved to be redundant and that the large amount of data actually confused us and concealed important patterns. Consequently, sensors were turned off in a systematic way until we were able to accurately "observe" the task oriented behavior with the use of the smallest number of sensors.

Stage 2: In this phase of testing ELMS was installed sequentially into the homes of volunteers. The installations ranged from one to three months and during this period, volunteers were interviewed on a regular basis to determine the validity of the recorded activities and the acceptability of the system. Perhaps the most important finding was that holidays and weekends changed people's activity pattern. In our first installation, we observed what appeared to be aberrant behavior for a single Thursday after three weeks of consistent activity. After pondering whether this was a problem with the hardware or software, or whether ELMS had pinpointed a previous unrecognized health problem, we interviewed the volunteer about his activity on this day. He replied, somewhat dumbfounded, "it was Halloween". Duh, people behave differently during holidays; a finding that has been confirmed over and over again since this revelation, but a finding that has required adjustments in algorithms to avoid false positives resulting in emergency alerts.

A second finding was that the volunteers "played" with the system for the first few days—they opened drawers and the refrigerator, removed medications frequently and waved hands in front of the sensors. When asked why, most volunteers replied that they were just "curious" and wanted to make sure the system was working by seeing the sensors flash. This "play" period lasted no more than two to three days after which the volunteers accepted the system and "forgot" that it was present. However, the increased "activity" affected the running average and we realized that it was not until at least the second or third week that "normal" activity would actually be determined.

Stage 3: ELMS was next installed in the residences of five elderly individuals in a high-rise, independent living senior building of one and two bedroom apartments in New York City. Although a social worker had access to a web site on which information was displayed, there was limited expectation that ELMS would be an integral part of caregiving. Once again, much noise emerged and much was learned. The first thing learned was that the thresholds for overnight bathroom use and meal activity were too sensitive, thus, generating false positive alerts. In the laboratory and even in the individual sequential testing phase, the thresholds worked well, but with very active people living in the real world of New York City, they were just too sensitive. This problem was easily corrected by adjusting the sigma values used for overnight bathroom use and meal activity.

The second lesson learned was more subtle, in some ways even more important and much more of a challenge to correct than the adjustment of sigma values. Since alert thresholds are based on deviation from an expected value equal to the running average of the number of events for the last 30 days, the thresholds are being continuously adjusted over time. Thus, changes in the number of events for a specific behavior can be of long-term significance, but because the changes are so gradual no alerts may be generated. Using overnight bathroom activity once again, there can be a steady increase in the number of events over a given 30 day period without producing either a single alert as the following series of real data points for a monitored individual shows--2, 2, 3, 2, 3, 3, 3, 2, 4, 3, 3, 4, 4, 3, 5, 4, 4, 5, 4, 5, 5, 4, 5, 6, 6, 5, 6, 6, 5, 6. The system is obeying its own rules for alert generation; however the change in the number of events, even though too gradual to generate daily alerts, may be very significant for the well-being of the monitored individual. Hence the conundrum: if the thresholds are set to be sensitive to these gradual changes, too many false positives are generated; if they are not adjusted, no alerts are generated and potentially serious health problems could be missed.

The solution was to develop a sub-routine that charted the average number of actual events per day for each of the behaviors for a designated time period--one week, two weeks, monthly. The sub-routine automatically counts the number of events for each behavior, produces a daily average and creates a multiple time period graph for each behavior. Therefore, for the example above, the average for the previous thirty days could be 2.2 while the average for the illustrated 30 days is 4.1. The resulting graph would clearly illustrate the magnitude of the change in the number of bathroom events between the two time periods and as a result, allows for gradual and steady changes in behavior to be observed even when no daily alerts were generated.

Stage 4: After the field test, we were ready to test ELMS in an environment in which care would be impacted by the information that the system generated. Thus, over an 18 month period, we installed ELMS in the homes of 34 chronically ill individuals living in their own residences in and around Philadelphia, PA. Given the complexity of this environment, both geographically—people lived miles apart, and structurally—visiting nurses had to access and respond to the information as part of their normal work day, the noise and the resulting challenges appeared, at times, insurmountable. However, we persevered and learned many lessons, the two most illustrative of which were: the need to develop a procedure for regular maintenance of the hardware; and that the information had to be customized for each individual care provider.

The maintenance issue, in hindsight, appears obvious, but except for replacing batteries in the sensors every 12-18 months—the sensors would even tell us when the batteries were low—we did not anticipate the need for a procedure for maintaining the hardware. We had not had problems with the hardware in either the laboratory or the more limited real world environments. We were wrong as a combination of poverty, serious health problems and unexpected living conditions conspired to drive home the point that if something can go wrong, it will and therefore, a plan must be in place to correct it. A trained individual must be ready to respond when sensors or base stations stop working; if not, the information generated by ELMS will degrade to the point of uselessness. Two brief examples illustrate the unexpected type of problems/noise we encountered. In the residence of one client, the sensors failed one by one over a two week period, eventually generating invalid information on all activities monitored. Since there was no one "charged" with the responsibility to go to the individual's residence to determine the problem and fix ELMS, the system was essentially out of action for close to a month. When someone finally went to the residence he found that all sensors and even the base station were infested with cockroach feces. Likewise, after ELMS abruptly stopped working in a second residence, over three weeks elapsed before someone made a special trip to the residence to determine the reason—a cat had chewed through the base station's power cord—and repair the problem. After these and several other examples, a maintenance protocol was developed that will be included in all future installations.

Lastly, we discovered that care givers have their own preferences as to how they want to receive the information generated. We assumed that the methods that we developed in the laboratory and were vetted by our focus groups, would be appropriate for overworked visiting nurses who would call in on up to eight to ten clients a day. Once again, we were wrong. Our protocol called for email alerts to be sent to the nurses and that after receiving the alerts the nurses would check the client's web page for more detailed information. Our assumption was that after reviewing all the information, the nurse would then determine what action was necessary. This assumption proved naïve. Each nurse developed his/her own emic protocol and wanted the information presented in a way that conformed to this construct. One nurse did not want alerts sent for a decline in meal activity for a single day. In his opinion, this told him nothing and just wasted his time. He only wanted an alert for meals if there were three consecutive days of abnormally low meal activity. One nurse never checked the client's web page and instead developed a protocol that required him to make a phone call to each of his clients who had any red alerts while a different care provider found the web pages to be of great use and requested even more detailed information on certain activities, e.g., sleep patterns. This has led us to build an ability to customize the way that information is provided to care givers into the implementation process of all future installations. As one care provider told us, "if it isn't given to me in a way I can use it, I won't even look at it".

In conclusion, although the real world challenges we have encountered have made the process of getting ELMS into the market place harder and longer than we anticipated, we have come to embrace the noise. What we once saw as encroaching chaos, we now view as the norm. The real world is messy, it is filled with unanticipated traps that frustrate and irritate, but if one stays in the safety of the laboratory there is little or no chance of successfully moving a new technology into the marketplace.