Mobility Assessment Using Event-Related Responses

M.Pavel[1], A. Adami[1], M. Morris[2], J. Lundell[2], T.L. Hayes[1], H. Jimison[3], J.A. Kaye[4,1]
[1]Department of Biomedical Engineering, Oregon Health and Science University, Portland, OR, USA
[2]Intel Digital Health Group, Intel Corp., Portland, OR
[3]Department of Medical Informatics, Oregon Health and Science University, Portland, OR, USA
[4]Department of Neurology, Oregon Health and Science University, Portland, OR, USA

Abstract—**The continuous assessment of mobility and speed of processing is an important component underlying physical and cognitive functions . We propose a novel approach to measure mobility, e.g. speed of walking and possibly speed of processing by unobtrusive monitoring of elders response times to specific events. The particular application investigated is response times to telephone ring. A key idea put forth in this paper is that if the elders' location distribution is stable over time, response times can be used to assess the "instantaneous" speed of walking. The feasibility of this approach is illustrated using data collected in a study performed by Intel in homes of several subjects.**

I. INTRODUCTION

THE ability to move is one of the critical functions that underlies the quality of life of elders [1]. In addition, changes in some aspects of mobility have been shown to correlate with changes in cognitive function and can perhaps predict future cognitive decline [2] The assessment of mobility is, therefore, an important procedure for clinical evaluation as well as determining the direction of care.

Cognitive and physical functions including speed of processing and mobility are generally assessed by measuring behavioral characteristics in a variety of settings. At one extreme, the behaviors are well structured and constrained in terms of laboratory tests. The subjects are brought into the laboratory or examining rooms and asked to perform specific tasks. Traditionally, the tests are low tech, for example using stop watches to time tasks such as walking a set distance (timed walk test) or connecting numbered dots on a piece of paper (trail-making test). Other versions of these tests utilize computers to measure reaction times on a sequences of trials. At the other extreme, subjective assessments of function are made by caregivers or trained observers reporting upon the elder in their home environment.

Recently researchers at the Oregon Health & Science University began to measure speed of walking in the home during normal daily activities, by placing motion sensors in the residences of subjects and unobtrusively collecting continuous measures throughout the day [3]. The motion sensors were placed in specially selected locations with known spatial arrangements such as hallways. Frequent passage through these parts of the residence allowed collection of sufficient data to determine typical walking speeds. The results were commensurate with those measured in the laboratory.

A shortcoming of this approach to measuring speed of walking is the requirement of a sufficiently long hallway and for relatively constant desire and motivation to walk at the same pace. Additional variance is likely to arise from various tasks that motivate subjects' movements and possible motivation to walk at different speeds. One approach to the measurement would be to ask subjects to perform certain, well-defined tasks during the course of their normal activities. This would, of course be obtrusive and would not allow continuous monitoring.

An alternative approach, described in this paper, represents an attempt to take advantage of responses to known events. There are several classes of everyday life events that generally trigger specific movements, including walking. These include various alarms, e.g., alarm clocks, doorbell, and telephone ring. The timing of these triggers as well as the responses to them could be monitored and used to assess speed of responses. A key hypothesis addressed in this paper is whether response times alone, i.e., without the knowledge of the locations, could be used for the assessment of elders' response speed.

II. GENERAL APPROACH: THEORETICAL ISSUES

In this paper we consider one such naturally occurring event – the response to an incoming telephone call. Although the response to a telephone ring is likely to be modulated by any concurrent tasks performed by the elder at the time of the ring, the response is likely to be more uniform than any general unconstrained mobility.

Under ideal conditions, the system would estimate the location of the elder at the time of the onset of the first ring and then measure the time to pickup and the location of the pickup. The elder's mobility defined by the speed of walking would be then assessed by the distance and time it took to pickup the telephone. This approach to assessment would require fusion of information about location including motion sensors with the information collected by telephone pickup sensor.

Due to limitations of currently available sensing technology, accurate localization may be difficult to obtain,

Manuscript received December 24, 2005. This work was supported in part by a grant from NIA P30 AG024978 and by a grant form Intel Corporation entitled "Intelligent systems for the early detection of dementia".

.

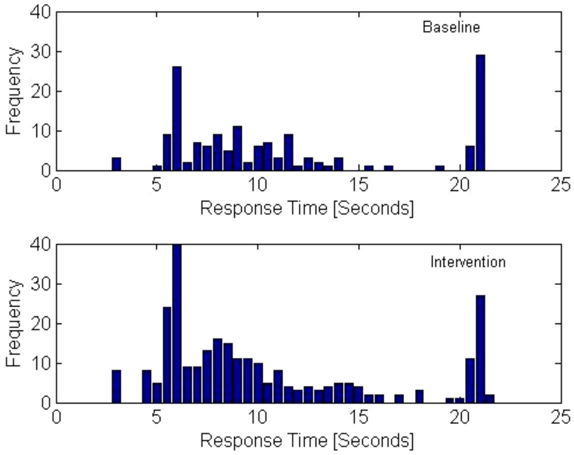

Fig. 1. An example of frequency distributions for response times of one subject during the baseline period (upper panel) and test interval (lower panel).

be uncertain, or be simply unavailable. For detecting changes in an individual's health state, the absolute value of the speed of walking may be less critical than its relative value, which allows assessment of change. We therefore examined the possibility of estimating these changes without the accurate knowledge of the elder's locations. This new approach is based on the assumption that the location of the elder at the time of the telephone ring is relatively stationary – i.e., that the distribution is approximately constant over time. This would be consistent with many elders' life styles – they may spend afternoons doing projects in the den and evenings watching television. In any case, this ergodicity assumption is empirically testable for each individual, as we shall demonstrate below.

If the assumption of ergodicity of spatial location is true, we can estimate the changes in speed of walking by examining the changes of the distribution of the response time to pickup the telephone over time, using the following theoretical framework. Let $f_D(d;t)$ denote the distribution of locations D (distances from the telephone set) over time t. If we assume that the walking speed v is a parameter that is relatively constant, i.e., v changes only slowly over time, then the response is a random variable $R = D/v$ with a distribution $f_R(r;t)$ which is derived from f_D by a simple scaling of the distance variable by a factor $1/v$. In that case, the changes of the distribution of the response times from time t_1 to t_2 can be described by a scaling factor, i.e., a multiplicative constant $\alpha = v_1/v_2$ representing the proportional change in the speed of walking from time t_1 to t_2. The procedure for estimation of a change in the speed of walking can therefore be reduced to the process of estimating the scale factor α.

If the distributions of the response times were well-behaved such that there were no outliers, then it would be possible to estimate the factor α. using ratios of sample moments, e.g., means and standard deviations,

$$\hat{\alpha} = w\frac{avg\langle R_2\rangle}{avg\langle R_1\rangle} + (1-w)\frac{SD\langle R_2\rangle}{SD\langle R_1\rangle}, \quad (1)$$

where w is relative weight and *avg* and *SD* are sample mean and standard deviation, respectively.

Since the ergodicity is not likely to hold exactly and since we wish to use small sample sizes, this approach is likely to lead to highly variable and biased estimates. The problem is in part due to the exaggerated influence of the samples from the distributions with higher expected values. We therefore use more elaborate approaches.

Prior to describing our approach to the estimation of α, we need to address the ergodicity assumption for the location probability distribution function. As we noted above, this is an empirical question and our answer is based on the analysis of the data of the Intel Social Health study. We therefore describe the Intel study before we focus on the details of the estimation process

III. INTEL SOCIAL HEALTH STUDY

The study carried out by Intel during the winter of 2004-2005 was designed to address two related aims:

(1) To examine the ability of a suite of sensors to infer the social health of an elder.

(2) To investigate whether various forms of visualization of the elder's social interaction would enhance his or her social interactions.

The data in the Intel study were obtained from ix households located in Las Vegas, NV. and Portland, OR. This study was aimed at monitoring and enhancing social interactions, but in this paper we focus on the subsystem involved only in data collection pertaining to interactions with the telephone. All subjects signed written informed consent to participate.

The telephone sensor was designed to detect most telephone activities and events, including off hook, on hook, ring, caller ID (whenever available for incoming calls), and DTMF dialing, including timing with millisecond accuracy. The study involved a number of other sensors, computer interfaces, and wireless networks, described elsewhere[4], that are not a part of the current data analysis. The study consisted of two 4-8 week intervals, baseline and intervention.

Any event-related information was time-stamped when it was received by the base station connected to the server PC. The data from each house was anonymized) with respect to identity and then used to populate an SQL database. The initial data cleanup and organizing were performed in the MATLAB© environment. For the purpose of this paper we focus on the measure of time-to-answer in response to a telephone ring.

A typical example of the response time distributions is shown in Figure 1. The upper panel shows the distribution during the baseline interval and the bottom panel shows the same distribution during the intervention interval.

There are two aspects of these graphs that are important with respect to the analysis of mobility. First is the similarity of the distributions for the two time periods. This similarity, if it is statistically significant, suggests that the distributions of the response times and thus the underlying location distribution did not change drastically between the first and the second time interval. These data, therefore, would provide the evidence for the assumption of the ergodicity of the location distribution. The second aspect of interest is the multimodal characteristics of the two distributions. The multimodal nature suggests that there are several locations where the elder can be found at the time of the first ring. Such distributions can generally be modeled by mixtures of distributions, such as Gaussian mixture model (GMM). An additional advantage of this interpretation is that whenever it is possible to estimate the location associated with a particular mode as well as the location of the telephone set, one could determine the absolute speed of walking.

IV. RESULTS

The purpose of this section is to explore ways to assess the similarity of the empirical distributions under several assumptions.

A. Stationarity assumption

As discussed above, the stationarity of the location distribution in conjunction with the assumption that response speed is the only random factor changing relatively slowly (on the scale of months), imply that any two distributions taken over non-overlapping time intervals should differ only in a scale factor α To examine whether this assumption holds we evaluated the response time distributions from baseline and intervention periods for each subject. We also estimated the scale factor relating each pair of distributions.

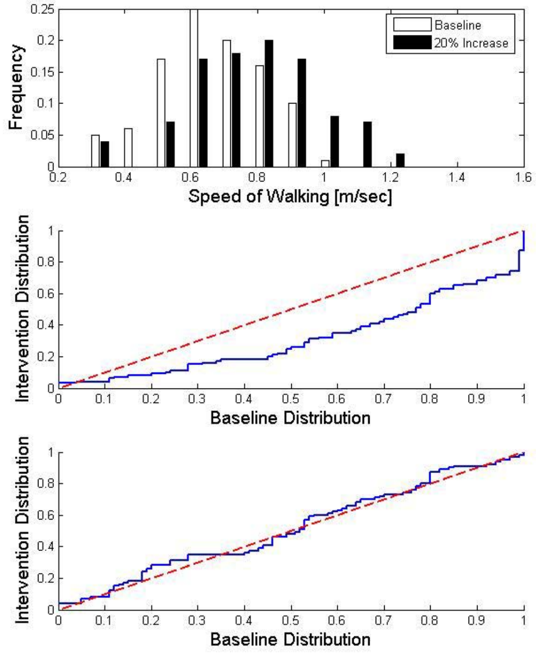

Fig 2. Simulation demonstrating two distributions and the corresponding stochastic dominance graph (ROC) that differ only in a scale factor. The lowest plate represents the same ROC after scaling by a factor $1/\alpha$

We first ascertain whether the shape of the distributions remains approximately the same. Informally, we explored this notion by a visual examination of the quantile-quantile (q-q) plots. The advantage of the q-q plots is that they are invariant with respect to affine transformations (shift and scale). The visual inspection supported the ergodicity assumption, except at the tails of the distributions where the variability is high due to small number of samples.

B. Pairwise assessment of the response time distributions

A quantitative approach to the comparison of distributions requires a choice of a suitable metric that is depends on the underlying assumptions. The classical statistical approach based on Kolmogorov-Smirnov test is not appropriate here because both distributions are empirical. Our approach is to combine a characterization of the relationship between the two distributions using stochastic dominance graph (see below) and then use an information-theoretic approach based on mutual information, or alternatively on Kullback-Leibler divergence, to quantify the difference between the baseline and intervention distributions.

In order to visualize the relationship between the two distributions that depend on the scale and shift parameters, we use a stochastic dominance graph. This graph is similar to the well-known receiver operating characteristic (ROC) and captures the probability that a variable chosen at random from one distribution is smaller (or greater for ROC) than one sampled from the other distribution. In order to illustrate the approach we simulated the problem at hand using a simple Gaussian distribution. We first sampled 100 points from a standard Gaussian distribution (baseline) with mean 9 and standard deviation of 2.2. We then sampled from the same distribution, but all the samples were scaled by the factor of 1.2, representing a 20% increase in response time (clinically meaningful change). The empirical histograms are shown in the top plate of Figure 2.

The “intervention” cumulative probability distribution function of the scaled variable is plotted against the “baseline” distribution function in the middle plate of Figure 2. If these two distributions were identical, all the points would lie on the diagonal line. The departure from the diagonal represents the differences in the distributions. The difference in the distributions can be quantified in a number of different ways, e.g., the Kolmogorov-Smirnov statistics or by the the area under the curve.

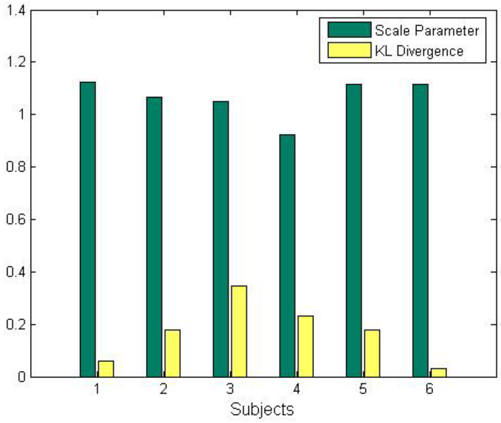

Fig. 4. Stationarity of the 6 subjects distributions as assessed by the KL-Divergence and estimate of the scale factor.

To illustrate the application of the stochastic dominance curve to the distributions of response time, we plotted these curves in Figure 3 for the same subject whose data was shown in Figure 1.

A useful metric for the difference between two distributions is the Kullback-Leiber divergence defined as follows

$$(p \| q) = E\left\{ p \log \frac{p}{q} \right\} \tag{2}$$

where p and q are two probability distribution functions. In case of discrete distributions this turns out to be the sum

$$D(p \| q) = \sum_{j=1}^{n} p(j) \log \frac{p(j)}{q(j)} \tag{3}$$

The procedure for estimating the parameter α (the proportional change in walking speed) can be implemented in terms of an optimization process than determines the value of α by minimizing the Kullback-Leiber divergence between the two distributions. This approach is generally preferable to the one based on moments, equation (1) because of the multimodal nature of the empirical distributions.

The values of the two parameters for each subject are compared in Figure 4. The fact that the value of the scale parameter is hovering around unity supports the notion that the speed of walking did not change significantly between the baseline and the intervention periods. An interesting conclusion from this comparison is that the KL Divergence is clearly measuring different aspects of the differences between the two distributions than is the scale parameter.

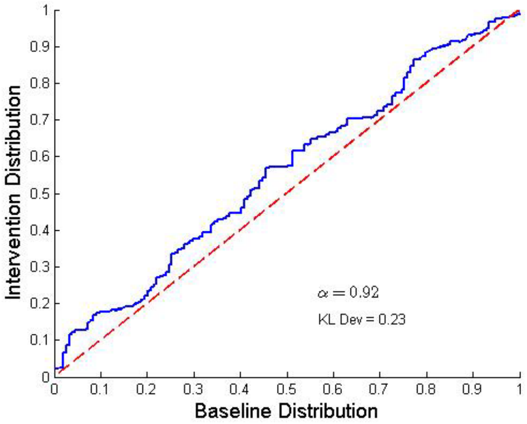

Fig. 3. An example of empirical stochastic dominance curve for one subject, for the distributions shown in Figure 1.

V. CONCLUSION

In this paper we discussed two novel ideas. First, we proposed an unobtrusive way to assess mobility and response speed in terms of response times to specific events during normal activities such as alarms, reminders or requests for responses. Second, we suggested a way to measure changes in the speed of walking that does not depend on measuring location of the elder at the time of the trigger event. We illustrated the general approach on a specific example involving the time that it takes elders to pickup a telephone call.

We explored the assumption that the distribution of location of the elder at the time of a telephone ring was stationary and speed of walking could be estimated by a scaling transformation. We note that the stationarity assumption is not required in those situations where the elder's location is assessed independently by other sensor..

ACKNOWLEDGMENT

We would like to thank our colleagues at Intel, namely Brad Needham and Terry Dishongh for their technical support and useful suggestions.

REFERENCES

[1] C. Callahan, H. Hendrie, and W. Tierney, "Documentation and evaluation of cognitive impairment in elderly primary care patients," *American College of Physicians*, vol. 122, pp. 422-429, 1995.

[2] R. Camicioli, D. Howieson, B. Oken, G. Sexton, and J. Kaye, "Motor slowing precedes cognitive impairment in the oldest old," *Neurology*, vol. 50, pp. 1496-8, 1998.

[3] T.L. Hayes, M. Pavel, J.A. Kaye,. "An Unobtrusive In-home Monitoring System for Detection of Key Motor Changes Preceding Cognitive Decline", 26th Annual International Conference of the IEEE Engineering In Medicine And Biology Society, 1-5 September, 2004; San Francisco, CA.

[4] M. Morris, "Social networks as health feedback displays", Internet Computing, 29-37, September 2005.