Postprint

This is the accepted version of a paper presented at *ACM SIGMOBILE International Workshop on Systems and Networking Support for Healthcare and Assisted Living Environments.*

N.B. When citing this work, cite the original published paper.

Permanent link to this version:
http://urn.kb.se/resolve?urn=urn:nbn:se:kth:diva-184702

Automatic Administration of the Get Up and Go Test

Dounia Berrada[1], Mario Romero[1], Gregory Abowd[1], Marion Blount[2], John Davis[2]

[1]College of Computing and GVU Center - Georgia Institute of Technology
{dounia, mromero, abowd}@cc.gatech.edu
[2]IBM T.J. Watson Research Center - Hawthorne, NY 10532
{mlblount, davisjs}@us.ibm.com

ABSTRACT

In-home monitoring using sensors has the potential to improve the life of elderly and chronically ill persons, assist their family and friends in supervising their status, and provide early warning signs to the person's clinicians. The Get Up and Go test is a clinical test used to assess the balance and gait of a patient. We propose a way to automatically apply an abbreviated version of this test to patients in their residence using video data without body-worn sensors or markers.

1. INTRODUCTION

The *Get Up and Go* (GUNG) test is used as a screening tool for the assessment of balance problems [Robertson, Mathias] that primarily targets elderly patients. Several variations of the test exist and they generally involve a patient starting in a seated position, standing, walking forward about 3 meters and then returning to the chair and sitting down. A person who performs the test in less than 10 seconds is assumed to be in good condition, while a person who performs the test in over 30 seconds is at risk for falling.

One of the problems with the GUNG test is that it is usually administered by a clinician or therapist in a clinical setting. We are investigating a system for automatically administering a variant of the GUNG test that will not require a medical expert. We feel that it is unrealistic to attempt to monitor the traditional GUNG test in a natural setting for several reasons. Mainly, it is unnatural. Hence, we are focusing on an abbreviated version of GUNG – Get Up and Go with First Step (GUNGFS). Using an in-home sensing infrastructure, we will deliver privacy enhanced GUNGFS test results to care providers for remote assessment.

Based on consultations with the medical community, we do not view the GUNGFS as a direct proxy for the standard GUNG test. Nevertheless, ongoing in-home monitoring of the GUNGFS will provide longitudinal data indicating patient trends over time that are simply not available to the medical community at present.

In this paper we present our research plan for designing, prototyping and experimenting with the automatically administered GUNGFS system along with early stage results. The two design points in our approach are low cost and easy setup. These design points drive a solution based on inexpensive, commodity hardware that can be setup in a home with little to no professional expertise. In the remainder of this paper we describe our system setup (Section 2) and then declare the research challenges that must be overcome to achieve our goal (Section 3). In Section 4 we present related work followed by status and conclusions in Section 5.

2. INITIAL SETUP

Our task is to automatically identify and time the activity of standing from a seated position and taking N initial steps within a home environment. We focus our work on activity monitoring through video capture based on two reasons. First, video capture does not burden subjects with wearing sensors. Second, video data enables health care providers to closely analyze interesting episodes.

2.1 Test setup

For our initial test environment we chose the living room couch which offers the advantage of having a relatively fixed position and is frequently used in most households and the disadvantage of often being a high traffic area. We considered three cameras (side, overhead and backrest) and selected the side camera because it produces footage that is most easy to review by humans observers. In our experiments, the distance from the side camera to the seat is 120 cm with a viewable standing height from the floor to 196 cm. The side camera we use is a Logitech QuickCam Pro 4000 that costs about $20. We fitted the camera with a wide angle lens with a field of view of 120° that retails for about $40.

2.2 Algorithms

Using the setup described above, we applied several standard image processing techniques to our collected video segments of subjects performing the GUNGFS and we produced encouraging results. For each video segment, we performed background subtraction using a fixed background template – an optimistic assumption that we consider to be reasonable. Upon subtraction and thresholding, all pixels containing motion (w.r.t. the background image) are referred to as active pixels and are white versus black pixels in which there is no motion. We reduce noise with the morphological operators *open* and *close*. Open gets rid of dusty looking active pixels assumed to be false positives. Close gets rid of holes in active pixel areas, assumed to be false negatives.

From the blobs of active pixels that remain we extract three features in both the vertical and horizontal axes: extent, mean and standard deviation as described below.

- *Extent* – The distance between the highest and lowest (rightmost and leftmost) active pixels in the vertical (horizontal) axis.

- *Mean* - The mean horizontal (vertical) value of the blob of active pixels.

- *Standard Deviation* – The statistical variance of the horizontal (vertical) values of active pixels.

Statistical aggregates such as the mean and the standard deviation (which take thousands of points to compute) are more robust to noise than single point measurements.

Figure 2 shows a graph of the vertical mean and standard deviation. The places where there are no values are the segments of the video where there are no foreground objects. Label A in Figure 2 points to an counterintuitive dip in the vertical mean that occurs as the subject stands up. When the subject is sitting (but starting to move in preparation for standing) movement occurs primarily in the upper torso and arms and there is very limited movement in the subject's legs. Hence, the vertical mean of the active pixels includes very few pixels below the seat of the chair. Once the subject stands, movement occurs not only in the torso but in the legs (at a height below that of the seat of the chair).

Using the features just described, we first detect when the person is standing. From this point we go backward until the torso movement on the couch before the standing sequence (we chose the first torso movement relative to a certain timeframe). Indeed we look for the first torso movement in the last 2 minutes before the standing position and we put a flag on this frame. When the person has moved one step away from the sitting position we place a "standing" flag. The time difference between these two frames is potentially the time it took the person to stand from the initial point of execution to the point of completion.

Figures 3 presents the vertical standard deviation for a subject performing a sequence of tasks. The tasks are: approach the couch, sit, attempt to stand, sit, stand, walk away, walk back, stand for a brief period, walk away, walk back, sit, attempt to stand, attempt to get up, stand, walk away, walk back, sit, stand, walk away, walk back, sit and stand with great effort. In our preliminary comparison of the data in Figure 3 with similar data from other subjects, we found significant correlation suggesting that our detection results are easily repeatable. We expect that more complex features or combinations of simple features will yield robust results.
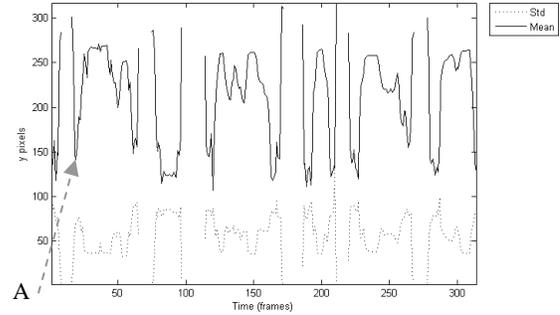

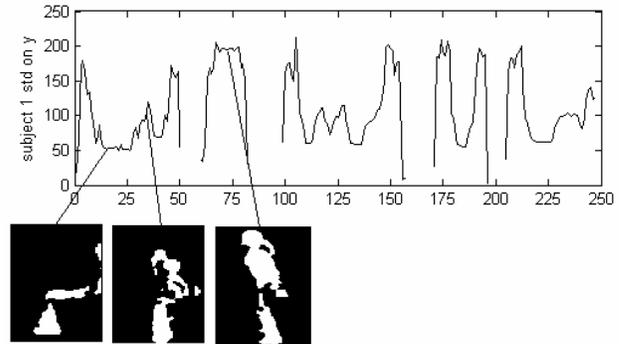Figure 2: Mean and Standard Deviation of Motion Pixels

Figure 3: Side Camera Results

3. OUR RESEARCH AGENDA

We see our initial test setup as a base from which to explore several research directions: easy deployment, resilient data processing and privacy. Below we outline these and related challenges.

3.1 System Deployment and Configuration

We want the patient (or a supporting family member) to be able to purchase a GUNGFS kit, bring it home, connect it to the Internet and "point" the corresponding camera at the patient's favorite chair. There is research in the area of no-configuration deployment of sensors that can be tapped to address the problem.

Once the initial setup is done, situations are not static. There may be intentional (rearrange furniture) or unintentional (camera was bumped) changes in the configuration. Our goal is to track the camera configuration so that, as needed, the patient can be guided in the correct configuration.

3.2 Noise Resistant Algorithms

The robust derivation of timing data from video is very challenging; interfering video noise or occlusions can directly impact the precision of our timing detection. We are considering the use of additional sensors to provide us supplemental information. For example, equipping the favorite chair with a padded pressure sensor can aid our system in focusing on episodic segments in which

GUNGFS is likely to be observed. Determination of the appropriate number of supplemental (non-worn) sensors will be key to avoiding over-learning from irrelevant sensors data. Monitoring multiple patients in a single home will require using biometric data for patient identification.

3.3 On Privacy

Based on feedback from the research literature and conversations with practitioners, we have to prepare for major resistance to the introduction of video capture in the home. One approach we are considering is to place a translucent blur lens in front of the video camera. Another approach we are considering is to disclose the video being recorded and receive patient approval before the data is sent to the care providers. We will investigate giving the person the ability to easily turn the system on and off. We also have to address the privacy concerns of visitors and other residents who are not the subject of the monitoring.

4. RELATED WORK

Several research projects have engaged in activity detection in a sensor rich home setting. The Aware Home Research Initiative [Kidd] at Georgia Tech and the MIT PlaceLab project [Intille] are examples. Researchers have explored activity detection based on sensors worn by persons. Lester et al [Lester], Charkson et al [Clarkson], [Taber] used variations of accelerometers, active badges and head mounted cameras. Other researchers like Niyogi & Adelson [Niyogi] and the UbiSense [Lo] project use cameras to monitor gait in patients. A key difference between much of the other video-based work and ours is that we attempt to accurately detect activity duration while timing is unimportant to the work of others.

5. STATUS AND CONCLUSIONS

We have reported on the initial phase of a research project that uses inexpensive video technology to make clinically significant assessments of the gait and balance of in-home patients. The results to date indicate there are paths to robust feature extraction and no-configuration deployment that are crucial in the long term.

6. ACKNOWLEDGMENTS

7. REFERENCES

[1] R.G. Robertson and M. Montagnini, "Geriatric Failure to Thrive," American Family Physician, 2004, 70:343-50.

[2] S.A. Niyogi and E.H. Adelson, "Analyzing and Recognizing Walking figures in XYT," Proc. IEEE Conf. Computer Vision and Pattern Recognition, pp. 469-474, 1994.

[3] S. Mathias, U. Nayak, B. Isaacs, Balance in the elderly patient: The "Get-up and Go" test. Archives of Physical Medicine and Rehabilitation, vol. 67, pp. 387-389, 1986.

[4] J. Lester, T. Choudhury and G. Borriello, "A Practical Approach to Recognizing Physical Activities," Proceedings of Pervasive 2006, LNCS 3968, pp 1-16, 2006.

[5] C. Kidd, R. Orr, G. Abowd, C. Atkeson, I. Essa, B. MacIntyre, E. Mynatt, T. Starner and W. Newstetter. "The Aware Home: A Living Laboratory for Ubiquitous Computing Research," In the Proceedings of the Second International Workshop on Cooperative Buildings, CoBuild'99.

[6] S. S. Intille, K. Larson, E. Munguia Tapia, J. Beaudin, P. Kaushik, J. Nawyn, and R. Rockinson, "Using a live-in laboratory for ubiquitous computing research," in Proceedings of PERVASIVE 2006, vol. LNCS 3968, K. P. Fishkin, B. Schiele, P. Nixon, and A. Quigley, Eds. Berlin Heidelberg: Springer-Verlag, 2006, pp. 349-365.

[7] A. Taber, A. Kesharvarz, H. Aghajan, Smart home care network using sensor fusion and distributed vision-based reasoning. In Proceedings of the 4th ACM international Workshop on Video Surveillance and Sensor Networks, VSSN '06. ACM Press, New York, NY, 145-154.

[8] B. Clarkson, K. Mase and A. Pentland, "Recognizing User Context via Wearable Sensors," International Symposium on Wearable Computers, pp 69-76, 2000.

[9] B. P. Lo, J. L. Wang, and G.-Z.Yang, "From imaging networks to behavior profiling: Ubiquitous sensing for managed homecare of the elderly." In Adjunct Proceedings of the 3rd Int'l Conf. on Pervasive Computing May 2005.