

Evaluation of an Inexpensive Depth Camera for Passive In-Home Fall Risk Assessment

Erik E. Stone

Department of Electrical and Computer Engineering
University of Missouri
Columbia, MO, USA
ees6c6@mizzou.edu

Marjorie Skubic

Department of Electrical and Computer Engineering
University of Missouri
Columbia, MO, USA
skubicm@missouri.edu

Abstract— We present an investigation of a new, inexpensive depth camera device, the Microsoft Kinect, for passive fall risk assessment in home environments. In order to allow older adults to safely continue living in independent settings as they age, the ability to assess their risk of falling, along with detecting the early onset of illness and functional decline, is essential. Daily measurements of temporal and spatial gait parameters would greatly facilitate such an assessment. Ideally, these measurements would be obtained passively, in normal daily activity, without the need for wearable devices or expensive equipment. In this work, we evaluate the use of the inexpensive Microsoft Kinect for obtaining measurements of temporal and spatial gait parameters as compared to an existing web-camera based system, along with a Vicon motion capture system for ground truth. We describe our techniques for extracting gait parameters from the Kinect data, as well as the advantages of the Kinect over the web-camera based system for passive, in-home fall risk assessment.

I. INTRODUCTION

To allow older adults to continue living longer in independent settings, and thus reduce the need for expensive care facilities, low-cost systems are needed to detect not only adverse events such as falls, but to assess the risk of such events, in addition to the early onset of illness and functional decline. Continuous, ongoing assessments of physical function would help older adults live more safely in independent settings, while also facilitating targeted medical interventions when needed. Ideally, such measurements would be obtained passively, in the course of normal daily activity [1].

This work focuses on developing a robust, low-cost, vision based monitoring system for assessing fall risk, detecting falls, and detecting the early onset of illness and functional decline. Research has shown the importance of measuring a person's gait, including identifying stride-to-stride variability as a predictor of falls [2-4]. Vision based monitoring systems have the resolution needed to yield the detailed measurements of physical function necessary for fall risk assessment (and early illness detection) passively, in the home environment, on a continuous basis. Furthermore, research has shown that the privacy concerns of older adults to video based monitoring systems can be alleviated through appropriate handling and processing of the video data, e.g., in the form of silhouettes [5].

Recently, Microsoft has released a new, inexpensive device, called the Kinect, to allow controller free game play on their Xbox system. The device uses a pattern of actively emitted infrared light to produce a depth image (the value of each pixel depends on the distance of what is being viewed from the device) which is invariant to visible lighting; and, thus, allows for a 3D representation using a single Kinect. This technology offers a number of potential benefits for low-cost, vision based monitoring systems.

This paper presents an investigation of the Kinect as a sensor for fall risk assessment. Specifically, techniques for acquiring spatial and temporal gait parameters from the depth data of the Kinect are presented; along with a comparison between the measurements obtained from the Kinect, to those obtained from an existing web-camera based system, and a Vicon marker based motion capture system.

II. BACKGROUND

Recent research in activity monitoring of older adults has focused on the use of passive infrared (PIR) motion sensor suites in the home [6-7]. These sensor suites yield information about the daily activity levels of monitored subjects, and arrays of such sensors have been used to obtain velocity measurements on a continuous basis in home settings [8]. While such systems don't raise privacy concerns among older adults, they typically do not produce measurements of the detail necessary for assessment of fall risk; specifically, spatial and temporal gait parameters (other than walking speed), timed up and go (TUG) time, sit to stand time, etc [9]. Existing systems for capturing such measurements are typically wearable, accelerometer based devices, expensive gait or motion capture systems, or direct assessment by a health care professional [2].

Wearable accelerometer based devices for obtaining detailed measurements of physical activity, specifically gait parameters, is an area that has been the focus of much research [10]. Efforts have even included utilizing accelerometers in existing smart devices, which individual's may already own and potentially carry with them. However, many elderly adults are reluctant to use wearable devices because they consider them to be invasive or inconvenient [1]. Thus, wearable devices may not be reliable for capturing movement in the home for continuous monitoring and

assessment.

Human motion analysis using video is a widely researched area, with two basic approaches: marker and marker-less. Marker based systems detect markers attached to a subject's body in multiple camera views. Given the location of the markers in different camera views, the 3D position of the marker can be obtained. The use of markers helps to yield highly accurate and robust measurements of a person's motion; however, marker based motion capture systems are not practical for in-home, continuous monitoring.

Marker-less video based motion capture systems generally work by extracting the silhouette of the subject in multiple camera views, and fitting a skeletal model to the intersection formed by the projection of the silhouettes in volume space [11-12]. Such systems have been shown to yield good results. However, they are typically expensive, require a controlled environment, or require high quality and/or a large number of cameras; attributes which limit their suitability for in-home activity monitoring. A number of researchers have looked at using systems composed of one, or small number of cameras for the purpose of detecting falls and monitoring activity in dynamic, home environments [13-20].

III. SYSTEMS OVERVIEW

A. Vicon

The Vicon system is a highly accurate marker based motion capture system used in a variety of animation, life sciences, and engineering applications [21]. The system outputs three-dimensional coordinates of detected markers at 100 frames per second. For this work, it serves to provide

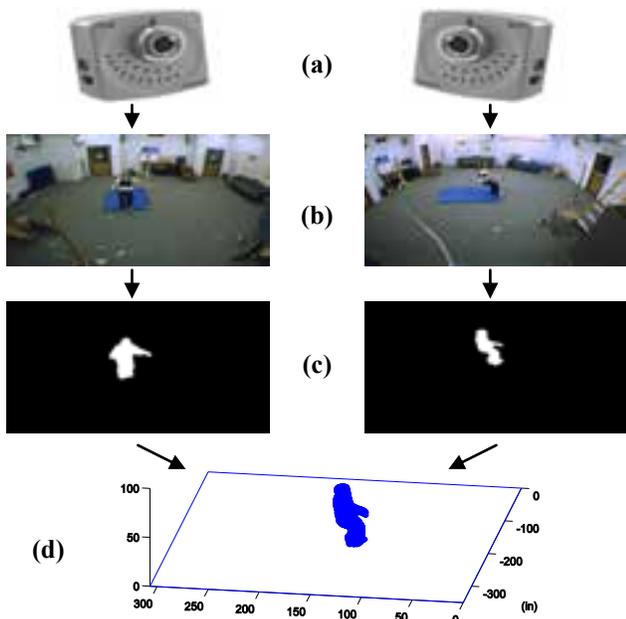


Fig. 1. Existing web-camera based system. (a) Cameras positioned orthogonally. (b) Views of the scene. (c) Extracted silhouettes. (d) Three-dimensional representation formed from silhouette projections.

ground truth data for comparison purposes.

B. Web-Camera

Our existing web-camera based system, outlined in Figure 1, consists of two inexpensive web-cameras, positioned roughly orthogonal, monitoring the environment. Silhouettes are extracted from captured images using a background subtraction and updating technique described in [22]. Intrinsic and extrinsic calibration parameters for both cameras are obtained *a priori*, allowing for a three-dimensional representation to be formed in the discretized volume space from the intersection of the projection of the silhouettes. Typically, the space is discretized into one inch (2.54 cm) cubic elements (voxels), and the system runs at approximately five frames per second, limited mainly by the silhouette extraction algorithm. The system has been evaluated for fall detection, gait measurement, and body sway measurement [23-25] with good results. The goal is to develop a passive, in-home, low cost, activity monitoring system for elderly adults.

C. Kinect

The Kinect, Figure 2, released by Microsoft in North America on November 4, 2010, was designed to allow controller free game play on the Microsoft Xbox. The device makes use of technology developed by the Israeli company PrimeSense, and contains both an RGB camera, and an infrared (IR) sensitive camera, from which a depth image can be produced based on a pattern of projected infrared light. The depth data returned from the device (at 30 frames per second) is an 11-bit 640x480 image which is invariant to visible lighting. The precision of the returned depth values is dependent on the distance, with the precision decreasing from approximately two centimeters at two meters to approximately ten centimeters at six meters. The minimum range of the device is approximately one meter. For use with the Microsoft Xbox, it is recommended the user be approximately two meters from the device.

When used with the Xbox, the Kinect allows 3D motion tracking using a skeletal model, gesture recognition, facial recognition, and voice recognition. Following its release, Linux and windows drivers were developed, and the Kinect has been used for a variety of purposes from entertainment to robotics [26].



Fig. 2. Microsoft Kinect.

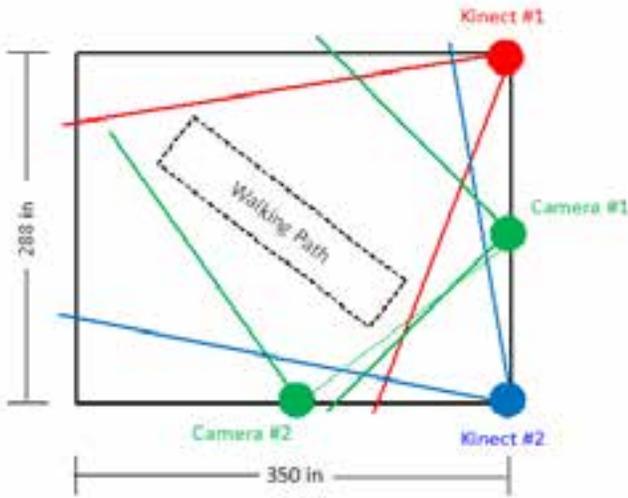


Fig. 3. Approximate position of web-cameras, Kinects, and walking path in test environment. Lines show field of view for each device.

D. Layout

For the experiments conducted in this work two web-cameras (forming our existing web-camera based system) and two Kinects were positioned in a laboratory environment, alongside a Vicon motion capture system. Figure 3 shows the placement of the web-cameras, Kinects, and the location of the walking path in the test environment. The cameras and Kinects were positioned approximately eight feet (2.4m) off the ground.

IV. METHODOLOGY

In this section, we describe the techniques used to extract the gait parameters of walking speed, right/left stride time, and right/left stride length from the 3D point cloud data returned by a single Kinect. For our purpose of evaluation, we make the assumption that there is only one person in the scene at a time, and that the environment is stationary. Thus, the 3D point cloud returned from the Kinect is for a single person. For use in a real world, dynamic environment, a high-level tracking algorithm would be necessary to achieve this. Finally, the extrinsic parameters of the Kinect with respect to the room have been computed. For the web-camera based system, existing algorithms were used to extract gait parameters [24].

A. Kinect – Calibration

The first step of obtaining accurate spatial parameters from the Kinect is calibration. This consists of two steps. First, as with traditional cameras, intrinsic, distortion, and stereo parameters for the IR and RGB cameras on the Kinect are estimated using a standard checkerboard calibration pattern and supplemental IR backlighting.

Second, calibration of the depth values returned from the Kinect is performed. The depth data returned from the Kinect must be transformed to obtain usable and accurate distances. For this work, the following equation, based on

[27], was used to transform the raw Kinect depth values, D , for a given pixel, (x, y) , to distances, d :

$$d = \frac{b}{f - D'}$$

$$D' = D(1 + k_1 r + k_2 r^2) + k_3 x' + k_4 y'$$

where x' and y' are the normalized pixel coordinates (computed using the intrinsic and distortion parameters of the IR camera), and r is the radius calculated using x' and y' . The parameters b , f , k_1 , k_2 , k_3 , and k_4 are optimized over a large (~3,000) set of training points; and the equation attempts to adjust for distortion effects. The training points are obtained by placing a large checkerboard calibration pattern in the environment, while moving the Kinect over a large range of distances and viewing angles with respect to the pattern. Using the known intrinsic parameters of the IR camera, the position of the calibration pattern with respect to the camera in each frame can be estimated. Simultaneously, the values associated with the pattern in the depth image can be recorded. Following collection of the training data, optimization of the parameters in the above equation is performed using the CMA-ES algorithm [28].

B. Kinect – Foreground Extraction

Foreground extraction is performed on the raw depth images from the Kinect using a simple background subtraction algorithm. Specifically, a set of training images is captured over which the minimum and maximum values for each pixel are stored to form a background model.

For a new frame, each pixel is compared against the background model and those pixels which lie outside the range formed by the minimum and maximum by greater than one are considered foreground. Next, a simple block based filtering algorithm is run to reduce noise, and smoothing is applied to the depth values identified as foreground. Example foreground extractions are shown in Figure 4.

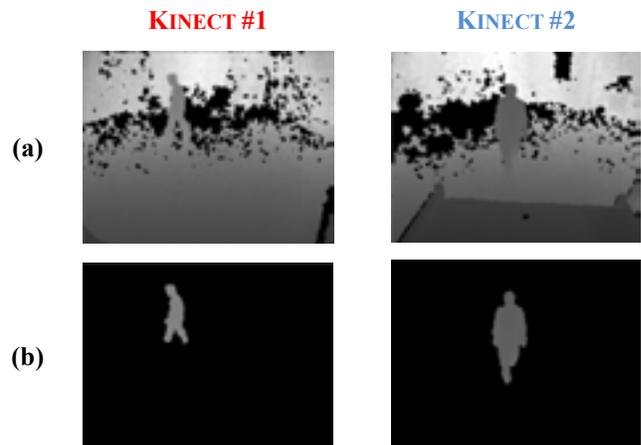


Fig. 4. Kinect foreground extraction. (a) Depth image from Kinect after distance conversion (black pixels mean no depth value was returned). (b) Extracted foreground.

This simple foreground extraction technique has proved to be quite robust, and easily runs at the 30 frames per second rate with which depth data is received from the Kinect. In a dynamic, real world setting, background adaptation to handle non-stationary environments would need to be addressed. However, the invariance of the depth image to changes in ambient lighting addresses one of the significant issues affecting foreground extraction on color imagery. Furthermore, the computation required for foreground extraction on the depth data is minimal compared to that required for robust foreground extraction from color imagery, where a combination of texture and color features must be used [29].

C. Kinect – Gait Parameters

Much of the work with the Kinect has focused on human body tracking using high degree of freedom skeletal models. Though such techniques are quite powerful, and may be essential to extracting certain physical parameters, they often suffer from problems of instability, especially with noisy data. As with our web-camera based system, we have opted to use techniques not based on skeletal models for extracting gait parameters.

Walking speed is estimated by projecting the centroid (smoothed using a moving average filter) of the 3D point cloud onto the ground plane, and measuring the change in position from frame-to-frame. For a given walking sequence, walking speed is computed by summing the frame-to-frame change in position over all the frames in the sequence, and dividing by elapsed time.

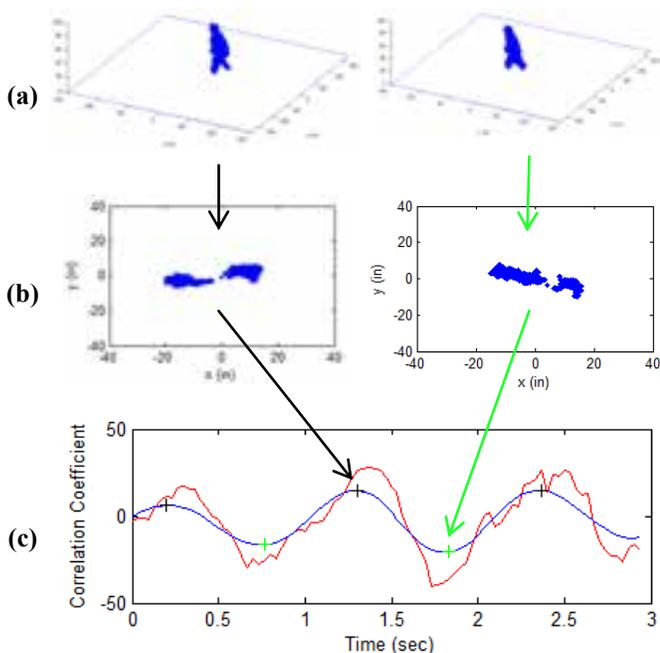


Fig. 5. (a) 3D point cloud from Kinect plotted in test environment. (b) Normalized ground plane projection of points below 20 inches. (c) Raw (red) and filtered (blue) correlation coefficient time series for walk sequence. Local maxima correspond to left footsteps, and local minima correspond to right footsteps.

The number of steps and temporal gait parameters are estimated using only those 3D points with a height below 20 inches (50 cm). In previous algorithms developed for the web-camera based system, only those voxel elements with a height of 4 inches (10 cm) or less were used for computing such parameters. However, due to the fact that the foreground extraction algorithm operates on the depth image, points from the person that are too close to the ground (and thus too similar to the background model) are not extracted as foreground. Furthermore, the distance at which a point is too close to the ground depends on the distance from the Kinect, as the measurement precision decreases as the distance from the Kinect increases. Therefore, points higher off the ground must be used.

First, such points are projected onto the ground plane. Second, the projection is normalized by subtracting the mean, and rotating based on the walking direction. Given the normalized projection, containing N points, the following correlation coefficient is computed:

$$\rho = \frac{\sum_{n=1}^N x_n y_n}{N}$$

where x_n and y_n correspond to the X and Y coordinates of the n^{th} point in the projection. The number of right and left steps for a walking sequence is obtained from the time series of the correlation coefficient for the sequence.

Specifically, the signal is first filtered using a moving average filter with a window size given by:

$$w = \frac{f * k}{v}$$

where v is walking speed, f is frame rate (for Kinect, 30 fps), and k is a constant parameter, although it could be adapted based on the estimated height of the person. The signal is then filtered a second time using a moving average filter with a small window size to remove any minor local extrema. From the filtered signal, right steps are detected as local minima, while left steps are detected as local maxima. Figure 5 shows example projections, along with a plot of the raw and filtered correlation coefficient time series for one walking sequence. The correlation coefficient of the normalized ground plane projection of 3D points below 20 inches has proven to be quite robust, even at large distance (over 20 ft.) from the Kinect.

Given the locations of the minima and maxima (right and left steps) in the correlation coefficient time series, the temporal gait parameters of right and left stride time can be computed. In addition, the spatial gait parameters of right/left stride length can be approximated as the distance moved by the centroid (as used in the computation of walking speed above) over the period corresponding to the right/left stride time. Although this approximation of the stride length may yield inaccurate measurements given large, abrupt changes in stride, it should still capture the

stride-to-stride variation which studies have shown to be predictive of falls [3-4].

V. RESULTS

We have collected gait measurement results from the different systems on a set of 18 walking sequences with varying subjects and speeds. The walking path is approximately 17 feet long, and the number of steps per walking sequence varies from five to nine. In half of the walks the subject was moving towards Kinect #2, and in the other half the subject was moving away. (Refer to Figure 3 for a placement diagram of the different sensor systems.)

A. Walking Speed

Figure 6 shows a plot of walking speed computed by each of the systems for each of the walking sequences, while Table I provides a comparison of the calculated percentage absolute difference (mean and standard deviation) between each system and the Vicon.

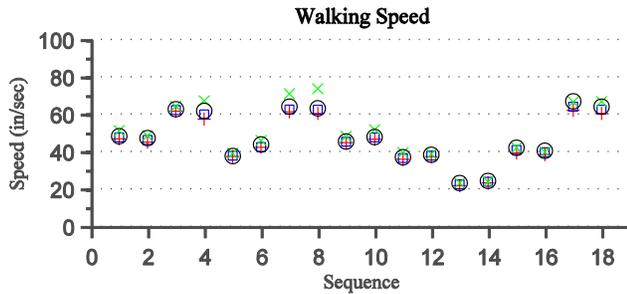


Fig. 6. Comparison of walking speed as measured for each of the 18 walking sequences. Vicon is black circle. Web-camera system is green x. Kinect #1 is red +. Kinect #2 is blue square.

TABLE I

% DIFFERENCE IN WALKING SPEED COMPARED TO VICON

	Kinect #1	Kinect #2	Web-Camera
Avg. % Diff	4.1	1.9	4.9
Std. Deviation.	1.9	1.1	4.3

As the results in Table I show, the individual Kinects yield good performance compared to the Vicon and web-camera system. Kinect #2 also seems to outperform Kinect #1 in terms of accuracy in walking speed measurement.

B. Stride Time

Figure 7 shows plots of the average right and left stride times, respectively, as computed by each of the systems for the 18 walking sequences. Table II provides a comparison of the calculated percentage absolute difference as compared to the Vicon.

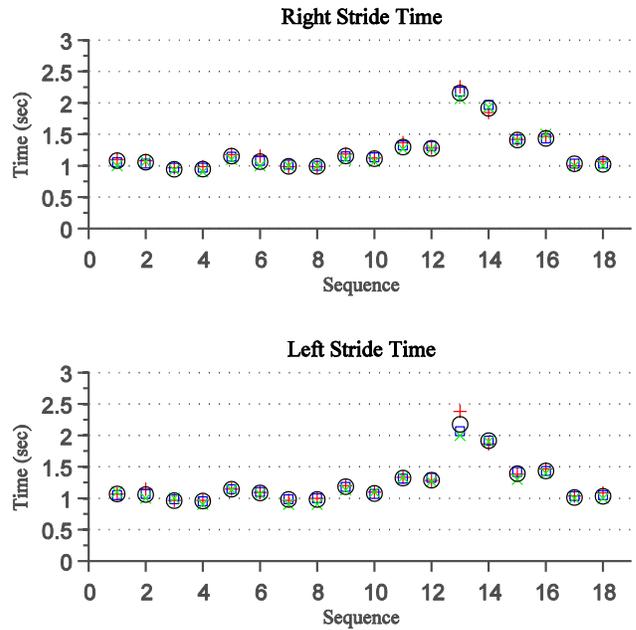


Fig. 7. Comparison of right and left stride time as measured for each of the 18 walking sequences. Vicon is black circle. Web-camera system is green x. Kinect #1 is red +. Kinect #2 is blue square.

TABLE II

% DIFFERENCE IN RIGHT STRIDE TIME COMPARED TO VICON

	Kinect #1	Kinect #2	Web-Camera
Avg. % Diff	2.8	1.6	3.0
Std. Deviation.	2.1	0.9	2.1

% DIFFERENCE IN LEFT STRIDE TIME COMPARED TO VICON

	Kinect #1	Kinect #2	Web-Camera
Avg. % Diff	2.3	1.4	3.8
Std. Deviation.	2.7	1.3	2.6

Interestingly, when it comes to measuring temporal gait parameters, the web-camera system should be at a significant disadvantage compared to the Kinects; due to its frame rate being six times slower. However, Table II seems to indicate the decreased frame rate does not cause a huge performance loss.

C. Stride Length

Figure 8 shows plots of the average right and left stride lengths, respectively, as computed by each of the systems for the 18 walking sequences. Table III provides a comparison of the calculated percentage absolute difference as compared to the Vicon.

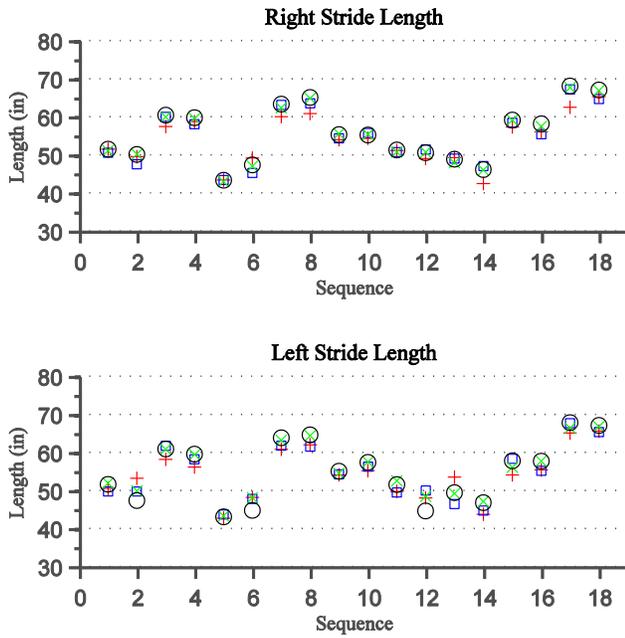


Fig. 8. Comparison of right and left stride length as measured for each of the 18 walking sequences. Vicon is black circle. Web-camera system is green x. Kinect #1 is red +. Kinect #2 is blue square.

TABLE III

% DIFFERENCE IN RIGHT STRIDE LENGTH COMPARED TO VICON			
	Kinect #1	Kinect #2	Web-Camera
Avg. % Diff	3.2	2.0	0.7
Std. Deviation.	2.6	1.6	0.5
% DIFFERENCE IN LEFT STRIDE LENGTH COMPARED TO VICON			
	Kinect #1	Kinect #2	Web-Camera
Avg. % Diff	5.0	3.7	1.9
Std. Deviation.	2.8	2.8	2.6

D. Discussion

Although this initial investigation of the Kinect for fall risk assessment showed both the accuracy of gait measurements made using the device and many potential benefits for fall risk assessment and in-home monitoring systems, there are issues that need further consideration.

First, certain types of clothing fail to reflect enough IR light back to the device to allow an estimate of depth at those pixels to be made. Furthermore, the issue of subjects blending into the background when they are close to walls, or, in the case of fall detection, on the ground is a concern in using the depth image alone for foreground segmentation and tracking. Potentially, a smart fusion of depth and color foreground segmentation could address some of these issues. Finally, another potential drawback of the Kinect is the limited field of view, approximately 60 degrees. This

restriction may require the use of multiple devices in many environments.

VI. CONCLUSION

As the Kinect has only been available for a short period of time, this work focused on evaluating the accuracy and feasibility of using the depth data obtained from the Kinect for passive fall risk assessment. Results showed good agreement between gait measurements computed using the Kinect, as compared to those computed using an existing web-camera based system, and those from a Vicon motion capture system. Furthermore, the depth image from the Kinect not only addresses a major issue in foreground extraction from color imagery (changing lighting conditions), but significantly reduces the computational requirements necessary for robust foreground extraction; potentially further reducing the cost of a fall risk assessment system.

Future work will look at obtaining and evaluating additional fall risk assessment parameters from the depth data of the Kinect, and will also explore a fusion of the depth and color imagery to achieve a fast, computationally inexpensive, and more robust foreground extraction than is possible with just the depth data or color imagery alone.

ACKNOWLEDGMENT

This work has been supported by the U.S. National Science Foundation under Grants IIS-0703692 and CNS-0931607.

REFERENCES

- [1] Demiris G, Rantz MJ, Aud MA, Marek KD, Tyrer HW, Skubic M & Hussam AA, "Older Adults' Attitudes Towards and Perceptions of 'Smarthome' Technologies: a Pilot Study," *Medical Informatics and The Internet in Medicine*, June, 2004, vol. 29, no. 2, pp. 87-94.
- [2] Hodgins, "The Importance of Measuring Human Gait," *Medical Device Technology*. 2008 Sep;19(5):42, 44-7.
- [3] BE. Maki, Gait changes in older adults: predictors of falls or indicators of fear, *Journal of the American Geriatrics Society*, vol. 45(3), pp. 313-20, 1997.
- [4] JM Hausdorff, DA Rios, HK Edelberg, Gait variability and fall risk in community-living older adults: a 1-year prospective study, *Arch Phys Med Rehabil* 2001;82:1050-6.
- [5] G. Demiris, O. D. Parker, J. Giger, M. Skubic, and M. Rantz, "Older adults' privacy considerations for vision based recognition methods of eldercare applications," *Technology and Health Care*, vol. 17, pp. 41-48, 2009.
- [6] TL. Hayes, M. Pavel, and JA. Kaye, "An unobtrusive in-home monitoring system for detection of key motor changes preceding cognitive decline," In: *26th Annual International Conference of the IEEE Engineering in Medicine and Biology Society*; San Francisco, CA, 2004.
- [7] S. Wang, M. Skubic, and Zhu Y, "Activity Density Map Dis-similarity Comparison for Eldercare Monitoring," Proceedings, *31st Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, Minneapolis, Minnesota, September 2-6, 2009, pp 7232-7235.
- [8] S. Hagler, D. Austin, T. Hayes, J. Kaye, and M. Pavel, "Unobtrusive and Ubiquitous In-Home Monitoring: A Methodology for Continuous Assessment of Gait Velocity in Elders," *IEEE Trans Biomed Eng*, 2009.
- [9] Guralnik, J. M., Simonsick, E. M., Ferrucci, L., Glynn, R. J., Berkman, L. F., Blazer, D. G., et al. (1994). A short physical performance battery assessing lower extremity function: association

- with self-reported disability and prediction of mortality and nursing home admission. *Journal of Gerontology*, 49(2), M85-94.
- [10] B. R. Greene, A. O'Donovan, R. Romero-Ortuno, L. Cogan, C. N. Scanaill, and R. A. Kenny, "Quantitative Falls Risk Assessment Using the Timed Up and Go Test," *IEEE Trans. on Biomedical Engineering*, vol. 57, no. 12, Dec., 2010, pp. 2918-2926.
- [11] TB. Moeslund, A. Hilton, V. Kruger, A survey of advances in vision-based human motion capture and analysis, *Computer Vision and Image Understanding (CVIU)* 104 (2-3) (2006) 90-126.
- [12] F. Caillette and T. Howard. Real-Time Markerless Human Body Tracking with Multi-View 3-D Voxel Reconstruction. In *Proc BMVC*. vol. 2, pp. 597-606, 2004.
- [13] I. Haritaoglu, D. Harwood, L.S. Davis, W4: real-time surveillance of people and their activities, in: *Proceedings of the IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 22, 2000, pp. 809-830.
- [14] C. Wren, A. Azarbayejani, T. Darrel, and A. Pentland, "Pfinder: Real-time tracking of the human body," In: *Proc. SPIE*, Bellingham, WA, 1995.
- [15] A. Mittal and L.S. Davis, "M2tracker: A multi-view approach to segmenting and tracking people in a cluttered scene using region-based stereo," In: *Proc. 7th European Conf. Computer Vision, Copenhagen, Denmark*, Vol. X, pages 18-33, 2002.
- [16] J. Krumm et al., "Multi-Camera Multi-Person Tracking for Easy Living," *Proc. 3rd IEEE Int'l Workshop Visual Surveillance*, IEEE Press, Piscataway, N.J., 2000, pp. 3-10.
- [17] H. Nait-Charif, S. McKenna, "Activity Summarization and Fall Detection in Home Environment," In: *Proc. of ICPR '04*, (2004) 323-326.
- [18] T. Lee and A. Mihailidis. An intelligent emergency response system: preliminary development and testing of automated fall detection." *Journal of Telemedicine and Telecare*, vol. 11, no. 4, 2005, pp. 194-198.
- [19] Z. Zhou, X. Chen, YC. Chung, Z. He, TX. Han, and JM. Keller, "Activity Analysis, Summarization, and Visualization for Indoor Human Activity Monitoring," *IEEE Transactions on Circuits and Systems for Video Technology*, 18:11, 2008.
- [20] B. Jansen, F. Temmermans and R. Deklerck, "3D human pose recognition for home monitoring of elderly," 29th Annual International Conference of the IEEE Engineering in Medicine and Biology Society, Lyon, France, Aug 23-26, 2007.
- [21] Motion Capture Systems from Vicon. <http://www.vicon.com/>
- [22] E. Stone, and M. Skubic, "Silhouette Classification Using Pixel and Voxel Features for Improved Elder Monitoring in Dynamic Environments," Accepted to *2011 IEEE International Conference on Pervasive Computing and Communications*, Seattle, WA, Mar 21-25, 2011.
- [23] D. Anderson, R. H. Luke, J. M. Keller, M. Skubic, M. Rantz, and M. Aud, "Linguistic Summarization of Video for Fall Detection Using Voxel Person and Fuzzy Logic," *Computer Vision and Image Understanding*, vol. 113, pp. 80-89, 2009.
- [24] E. Stone, D. Anderson, M. Skubic, and J. Keller, "Extracting Footfalls from Voxel Data," 32nd Annual International Conference of the IEEE Engineering in Medicine and Biology Society, Buenos Aires, Argentina, Aug 31-Sep 4, 2010.
- [25] F. Wang, M. Skubic, C. Abbott, and J. Keller, "Body Sway Measurement for Fall Risk Assessment Using Inexpensive Webcams," 32nd Annual International Conference of the IEEE Engineering in Medicine and Biology Society, Buenos Aires, Argentina, Aug 31-Sep 4, 2010.
- [26] OpenKinect. <http://openkinect.org>
- [27] Technical Description of Kinect Calibration. http://www.ros.org/wiki/kinect_calibration/technical
- [28] N. Hansen, "The CMA Evolution Strategy: A Comparing Review," In *Towards a new evolutionary computation. Advances in estimation of distribution algorithms*, pp. 75-102, Springer, 2006.
- [29] R. H. Luke, "Moving Object Segmentation from Video Using Fused Color and Texture Features in Indoor Environments", Technical Report, CIRL, University of Missouri, 2008.