Automated In-Home Assistive Monitoring with Privacy-Enhanced Video

Alex Edgcomb
Department of Computer Science and Engineering
University of California, Riverside
aedgcomb@cs.ucr.edu

Frank Vahid
Department of Computer Science and Engineering
University of California, Riverside
vahid@cs.ucr.edu

ABSTRACT

A privacy-enhanced video obscures the appearance of a person in the video. We consider four privacy enhancements: person blurred, person silhouetted, person covered with a bounding-oval, and person covered by a bounding-box. We demonstrate that privacy-enhanced video can be as accurate as raw video for eight in-home assistive monitoring goals: energy expenditure estimation, in room too long, leave but not return at night, arisen in morning, not arisen in morning, in region too long, abnormally inactive during day, and fall detection. Each monitoring goal's solution was trained using one actor and tested using two different actors. The privacy enhancements of silhouette, bounding-oval, and bounding-box, did not degrade achievement of the eight assistive monitoring goals. Raw video had a fidelity of 0.994 for the goal of energy expenditure estimation, while silhouette had 0.995, bounding-oval had 0.994, and bounding-box had 0.997. The fall detection algorithm yielded the same sensitivity of 0.91 and specificity of 0.92 for raw and bounding-oval video, while silhouette had a sensitivity of 0.91 and specificity of 0.75, and bounding-box had a sensitivity of 0.82 and specificity of 0.92. The other 6 goals yielded perfect sensitivity and specificity for raw and privacy-enhanced video, with the exception of blur video's sensitivity of 0.5 in region too long.

Categories and Subject Descriptors

J.3 [**Computer Applications**] - *health*.

General Terms

Algorithms, Measurement, Reliability, Experimentation, Human Factors.

Keywords

Assistive monitoring, smart homes, privacy-enhanced video, telehealth, embedded systems, ubiquitous systems.

1. Introduction

Automated monitoring algorithms operating on live video streamed from a home can effectively aid in several assistive monitoring goals, such as detecting possible falls or determining activity levels. Use of video raises obvious privacy concerns. Several privacy enhancements have been proposed such as modifying a person in a video by adding a blur, silhouette, or bounding-box. We investigate the impact of 4 common privacy enhancements on the effectiveness of 8 automated monitoring goals. We show that privacy-enhanced video can be as accurate as raw video for the goals of detecting: a fall, in room too long, leave but not return at night, arisen in morning, not arisen in morning, in region too long, and abnormally inactive during day. We also show that the goal of energy expenditure estimation is only 5% less accurate with privacy-enhanced video than with raw video.

2. Background

Raw video is quite-obviously perceived to have less privacy than privacy-enhanced video [8][12]. Demiris [8] surveyed 15 participants over the age of 65, all of whom felt that the use of a camera for in-home monitoring was obtrusive, while "many participants felt that [silhouetting] was more appropriate," (numerical data was not provided). Our earlier work [12] surveyed 328 participants (average age of 20 years) with raw, person covered by blur, person replaced by silhouette, person covered by filled oval, person covered by filled rectangle, and person replaced by blue-outlined rectangle with trailing arrows videos. The effectiveness of the privacy enhancement was rated on two metrics: privacy score and sufficiency of privacy. The privacy score was based on a 6-choice Likert scale (strongly disagree, disagree, slightly disagree, slightly agree, agree, strongly agree) for the statement, "This [privacy enhancement] protects Grandpa's privacy". The privacy score's range was 0 to 18, in which 0 meant strongly disagree and 18 meant strongly agree. Raw's privacy score was 2.4, whereas blur had 9.5, silhouette had 11.6, bounding-oval had 14.0, bounding-box had 15.5, and trailing-arrows had 16.0. The sufficiency of privacy was the percentage of participants who responded that a privacy enhancement provided a sufficient amount of privacy. Only 2% said that raw video provided sufficient privacy, whereas blur had 23%, silhouette had 59%, bounding-oval had 88%, bounding-box had 96%, and trailing-arrows had 98%.

Commercial assistive monitoring systems typically use only traditional sensors to detect anomalies and provide configurable event-of-interest detection. BeClose [2] monitors traditional sensors, such as motion sensors and door sensors, for anomalies in daily activities and notifies a caregiver in the event of an anomaly. Motorola's Homesight [18] is configured by the caregiver or a technician to detect

Figure 1: Our in-home environment included a living and dining room. (a) Picture taken from camera used for recording. (b) Floor plan of environment showing the position and direction of the camera used for recording. (Floor plan created with http://floorplanner.com).

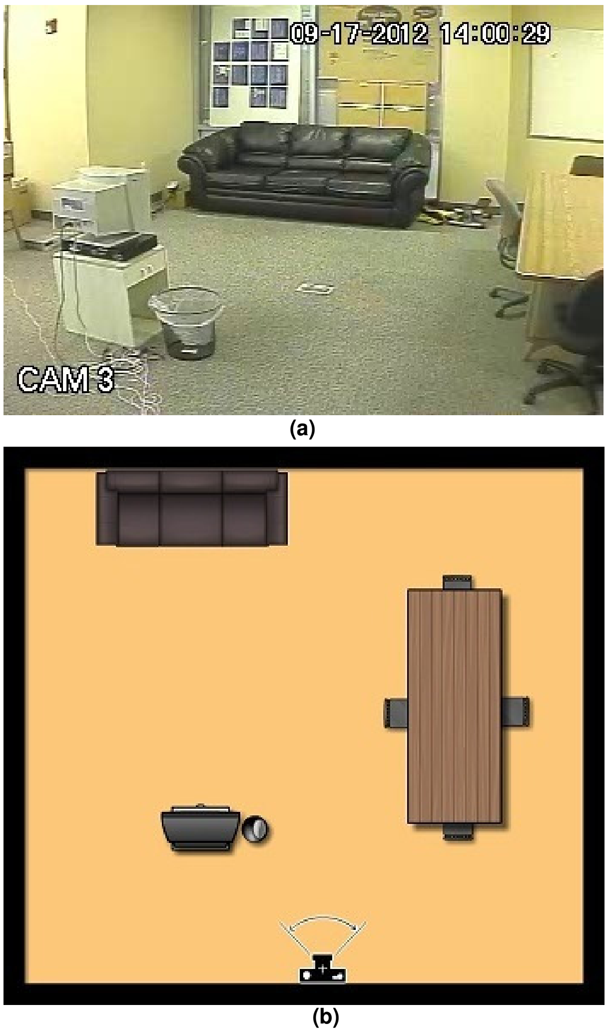

events of interest, such as a person leaving home at night. A camera can be configured to take a picture or record video when an event of interest is detected. QuietCare [20] performs anomaly detection after learning typical patterns with motion sensors that are placed throughout the home and notifies the caregiver when an anomaly is detected. SmartHome [23] and X10 [25] offer a variety of monitoring kits and products, such as an eight camera with basestation security system and a programmable thermostat.

Academic assistive monitoring systems typically use traditional sensors to detect anomalies and provide programmable platforms. CASAS [21] performs anomaly detection on sensors, such as motion and light sensors, by learning normal behavior then detecting salient events. The Gator Tech Smart House [14] is a sensor platform that is programmable for detecting events of interest or anomalies. The types of sensors are vast, including floor sensors, plug sensors, and smart microwaves; however, cameras are not used to monitor the interior of the home due to privacy concerns [13].

Automated fall detection has been performed with on-body accelerometers [3], and off-body sonar [17] and cameras [1][22]. The accelerometers and sonars provide high-levels of privacy in that the identity and specific activity is not discernible. Some camera-based solutions are privacy-enhanced, such as Anderson [1] that uses extracted person-silhouettes to detect falls.

3. Recordings

Our recording environment was a mock-apartment that included a living room and dining room, shown in Figure 1. The environment was located in a research laboratory at the University of California at Riverside. Each recording contained a single actor. In some recordings, the actor performed one or more general events, such as entering or exiting the apartment. In some recordings, the actor performed a specific task, such as read on couch or sweep the floors. The recordings were grouped by monitoring goal, and each group had a set video length, such as 30 minute recordings for in room too long detection.

The camera was apart of a popular 8-camera in-home monitoring set made by Q-See and sold by Costco for around \$400. The camera was 182cm above the floor (approximately head-level of the actors).

A 27-year old UCR graduate student researcher on this project and two UCR undergraduate student volunteers, ages 22 and 21. All male.

4. Foregrounding via foreground-background segmentation

A moving object in the foreground is extracted from a static background via foreground-background segmentation. For in-home assistive monitoring, the common moving object is a person, particularly in the home of a live-alone elderly person. Foregrounding is a typical problem in computer vision and many algorithms have been developed [15][16][24]. The main idea of the algorithms is that non-changing parts of the video are learned as a background image. The background image is subtracted from a video frame, resulting in only the foreground image, which contains moving objects.

We used the established foregrounding algorithm by KaewTraKulPong [15] to extract a foreground image, shown in Figure 2(c). Using the output foreground image, we developed an algorithm [10] to place a minimum bounding rectangle (MBR) around the largest object in the foreground image, as shown in Figure 2(d).

5. Privacy enhancements

Cameras typically output raw video, which we define as follows:

Figure 2: Foregrounding via foreground-background segmentation takes as input (a) a background image and (b) a video frame, and outputs (c) a foreground. (d) A minimum bounding rectangle (MBR) is fit around the foreground.

- *Raw video*, shown in Figure 3(a), is normal video that shows the camera's scene as clearly as possible.

A privacy-enhanced video intentionally obscures the appearance of a person in the video to protect that person's privacy. Raw video is perceived to have less privacy than privacy-enhanced video [8][12]. We consider four privacy enhancements:

- *Blur video* (Figure 3(b)) smears the video, typically restricting the smearing to the region with movement.
- *Silhouette video* (Figure 3(c)) covers the movement with an outline of the person filled with a solid color.
- *Bounding-oval video* (Figure 3(d)) covers the movement with a bounding oval around each person.
- *Bounding-box video* (Figure 3(e)) covers the movement with a bounding box around each person.

We built a tool to convert raw video to privacy-enhanced video. The raw video was processed with our foreground-background segmentation algorithm to extract a foreground and the MBR around the foreground. The blur video blurred the region of the raw video in which the MBR resides. The silhouette video changed a pixel to blue if that pixel was part of the foreground and within the region of the MBR. The bounding-oval video covered the region of the MBR with a solid blue oval that has the same height and width of the MBR. The bounding-box video covered the region of the MBR with a solid blue rectangle with the same height and width of the MBR.

6. Camera-based assistive monitoring goals

Live-alone persons, particularly elderly live-alone persons, may wish to be monitored for situations of interest that indicate a problem. The assistive monitoring goals would be determined by the monitored person or that person's caregiver, which could be an adult child or a nursing home staff. Some goals can be solved with a camera-based approach, such as: energy expenditure estimation, in room too long, leave but not return at night, arisen in morning, not arisen in morning, in region too long, abnormally inactive during day, and fall detection.

Figure 3: Pictures from the same moment of the same recording as (a) raw, (b) blur, (c) silhouette, (d) bounding-oval and (e) bounding-box.

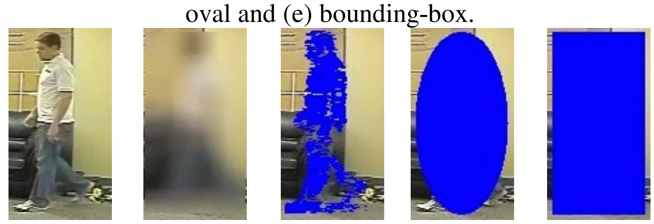

The camera could be installed by the monitored person, the caregiver, or a trained technician. Similarly, the assistive monitoring goals using the installed camera could be configured by the monitored person, caregiver, or a trained technician using an assistive monitoring language such as the Monitoring and Notification Flow Language [11].

6.1 Energy expenditure estimation

Automatically estimating a person's daily energy expenditure has numerous uses, including determining whether an elderly person living alone is achieving sufficient levels of daily activity. The U.S. Centers for Disease Control and Prevention and various researchers suggest moderately-intense physically activity for at least 15 minutes, 3 or more times a week [7][19]. Most previous automatic energy expenditure estimation approaches require use of body-worn devices [5][6][26]. Body-worn devices have advantages relating to their accuracy and to their measurement of activity inside or outside the home. However, such devices have some key drawbacks, including that people commonly fail to wear the device (either intentionally or unintentionally) [4].

In previous work, we developed an algorithm to estimate energy expenditure using raw video [10]. In this work, we evaluate the algorithm's effectiveness with the four privacy enhancements.

6.2 In room too long

A person inside a room, such as a study, laundry room, and bathroom, for an extended period of time may indicate a problem. One solution is to start a timer when a person enters a specified room. If the person does not leave before the timer runs out, then send a notification to the caregiver. A person can be detected as entering a room by positioning a camera perpendicular to a room's entrance such that a person enters the room by leaving the camera's view to the left.. Similarly, the person exits the room by entering the camera's view from the left.

We positioned a camera outside the room perpendicular to the entrance/exit of the room such that a person enters the room by leaving the camera's view to the left, and exits the room by entering the camera's view from the left. We built an in-room-too-long detector by starting a stopwatch when the MBR (representing the person) exited the camera's view to the left and resetting the stopwatch when the MBR entered the camera's view from the left. If the stopwatch exceeded 20 minutes, then the caregiver is notified.

6.3 Leave but not return at night

Particularly for Alzheimer's patients, a person who leaves their house at night but does not return within 15 minutes may indicate a problem. One solution is to start a timer when the person leaves the house between 10PM and 4AM. If the person does not return before the timer runs out, then send a notification to the caregiver. A person can be detected as leaving by positioning a camera perpendicular to the house's entrance/exit such that a person exits the house by leaving the camera's view to the left. Similarly, the person enters the house by entering the camera's view from the left.

We modified the in-room-too-long detector to create the leave-but-not-return detector by including an and-operation between the in-room-too-long detector and whether the current time was night. This modification prevents a notification from being sent when the person leaves during the day for an extended period of time.

6.4 Arisen and not arisen in morning

A basic concern of a caregiver for an elderly person is when the elderly person arose that morning. One solution is to send a notification to the caregiver when the first significant motion has been seen outside the bedroom between 6AM and 11AM. A person can be detected as arisen by positioning a camera in the main living area. The person's movement in the main living area will generate a significant amount of motion thus indicating the person has arisen. We positioned a camera to view the main living area. The first time between 6AM and 11AM that the MBR (representing a person) moves, a single notification is sent.

Another basic concern for caregivers is to detect that a person has not arisen by 10AM. One solution is to send a notification to the caregiver if no significant motion has been seen outside the bedroom between 6AM and 10AM. A person can be detected as having arisen by positioning a camera in the main living area. The person's movement in the main living area will generate a significant amount of motion.

6.5 In region too long

A person being in a particular region, such as a hallway, for an extended period of time may indicate a problem. One solution is send a notification to the caregiver if the person is inside a specified region for over 10 minutes. A person can be detected as being in a hallway by positioning a camera to view the length of the hallway.

We positioned a camera in the hallway viewing the length of the hallway. The person, represented by an MBR, is compared to the specified region to determine whether the MBR is within the region. We started a stopwatch when the person, represented by an MBR, was within the hallway region and reset the stopwatch when the MBR was not within the region. The caregiver would be notified if the stopwatch reached 10 minutes.

6.6 Abnormally inactive during day

A person having a particularly inactive day may indicate a problem. One solution is to notify the caregiver if the person is at home but has not moved for 3 hours. A person can be determined to be at home by positioning a camera perpendicular to the house's entrance/exit such that the person exits the house by leaving the camera's view to the left, and enters the house by entering the camera's view from the left. A person can be determined to not have moved for 3 hours by running a timer that is reset whenever a camera positioned in the main living area detects significant motion.

We positioned a camera perpendicular to the house's entrance/exit and viewing the main living area. A stopwatch started when the person was home but no significant motion was detected, and the stopwatch was reset when either the person left home or significant motion was detection. A notification was sent to the caregiver if the stopwatch reached 10 minutes.

6.7 Fall detection

A common goal in assistive monitoring is automated fall detection. In previous work, we developed a fall detection algorithm for raw and privacy-enhanced video [9]. A feature of the MBR was used to detect falls, in which we found that the width of the MBR in pixels to be accurate. The width of the MBR has a characteristic time series when a fall occurs. To detect falls, we compared an observed width of MBR's time series to a characteristic fall time series. The key observation from this work is that raw and privacy-enhanced video have similar MBRs for the same video recording.

7. Experiments of camera-based assistive monitoring goals

The monitoring environment may be dramatically different between homes and persons. Also, camera placement will modify the specific implementation of goal solutions. Therefore, we assume a training phase to improve the goal solutions. The training phase may be re-run periodically, as goals, monitoring environments, and monitored persons change over time.

The person being monitored may not have the physical capacity to train the monitoring system his or her self. Therefore, we assume that either the caregiver or a trained technician would train the solution for each assistive monitoring goal.

7.1 Energy expenditure estimation

A key expected use of video-based activity level estimation is to compare a person’s activity levels across many days, to detect negative trends and thus introduce interventions. As such, a goal of estimation is not necessarily accurate calorie estimation, but rather correct relative estimation of energy expenditure across days, including correct ratios among low/medium/high activity days. Thus, our experiment sought to determine the *fidelity* of our video-based activity level estimation. We define fidelity as the correlation between the video-based energy estimation and the BodyBugg energy estimation (such a definition tolerates an offset between estimated and actual calories). The ideal fidelity is a correlation of 1.0. However, the accuracy of the energy estimate is also potentially useful for determining whether a particular Calorie expenditure goal was met. We define accuracy as 1 minus the approximation error, which is the absolute difference between the expected value and the observed value divided by the expected value. Accuracy of 100% is ideal.

We conducted experiments to determine the fidelity and accuracy of privacy-enhanced video. We used our previous works' video-based energy expenditure estimation algorithm, video data set, and the same testing methodology, which included four actors and each actor having three 16-hour constructed days: a low activity level day, a medium activity level day, and a high activity level day [10].

Table 1 shows the average accuracy and fidelity of energy expenditure estimation with raw and privacy-enhanced video. Raw video had a higher average accuracy than all of the privacy enhancements. Bounding-oval had the highest average accuracy of the privacy enhancements but was still 5.3% behind raw video. The average fidelity was close to 1.0 for raw video and each privacy enhancement with bounding-box being the highest of all at 0.997.

7.2 In room too long

We conducted experiments to determine the sensitivity and specificity of raw and privacy-enhanced video of the in-room-too-long detector. Sensitivity is the ratio of correct in-room-too-long detections over the actual number of in-room-too-long situations. Specificity is the ratio of correct not in-room-too-long detections over the actual number of not in-room-too-long situations. The detector's enter/exit from left was trained by Actor 3 using only raw video in three situations, each 30 minutes long: person enter room but not leave in time (situation of interest), person enter room then leave room in time, and person nearly enter room. The detector was tested with Actor 1 and Actor 2 using raw and privacy-enhanced video in the same three situations. The average sensitivity and specificity for raw and each privacy-enhanced video were all 1.0. A perfect result is not surprising considering how simple detecting the direction is from which a person enters or leaves a camera's view.

Table 1: Average accuracy and fidelity of energy expenditure estimations with raw and privacy-enhanced video. Higher is better. Privacy-enhanced video had lower accuracy than raw video by 5.3% to 10.4%, while the fidelity for bounding-box was higher than raw video.

Privacy enhancement	Average accuracy	Average fidelity
Raw	90.9%	0.994
Blur	80.5%	0.991
Silhouette	85.0%	0.995
Bounding-oval	85.6%	0.994
Bounding-box	84.3%	0.997

7.3 Leave but not return at night

We did not conduct an experiment specifically for this assistive monitoring goal because the in-room-too-long detector was already evaluated and the modification was a minor logic change (a simple and-operation with current time). Instead, the average sensitivity and specificity are assumed to be the same as the in-room-too-long detector, which were both 1.0.

7.4 Arisen and not arisen in morning

We conducted experiments to determine the sensitivity and specificity of raw and privacy-enhanced video of the arisen-in-morning detector. The amount of pixels the MBR was required to move was trained by Actor 3 using only raw video in four situations, each 5 minutes long: person enter room from left side of camera's view, person enter room from right side of camera's view, person already in room, and person never enter room. The detector was tested with Actor 1 and Actor 2 using raw and privacy-enhanced video in the same four situations. The average sensitivity and specificity for raw and each privacy-enhanced video were all 1.0.

We modified the arisen-in-morning detector to send a message if no motion had been seen between 6AM and 11AM. We did not conduct an experiment specifically for this assistive monitoring goal because the arisen-in-morning detector was already evaluated and the modification was a minor logic change.

7.5 In region too long

We conducted experiments to determine the sensitivity and specificity of raw and privacy-enhanced video of the in-region-too-long detector. The specific rectangular region of pixels was trained by Actor 3 using only raw video in three situations, each 15 minutes long: person partway in region, person completely in region for 5 minutes then partway in region for 10 minutes, and person completely in region. The detector was tested with Actor 1 and Actor 2 using raw and privacy-enhanced video in the same three situations. The average sensitivity and specificity for raw and each privacy-enhanced video, except blur, were all 1.0. Blur video had a sensitivity of 0.5 and specificity of 1.0. The failed detection was from "person completely in region" for actor 1's blur video, in which the MBR slowly shrunk as pixels between the blurred person in the foreground and the background appeared extremely similar. The foreground-background segmentation algorithm uses the color difference in pixels between the foreground and background. Blur video reduces the difference between the foreground and background by blurring. Therefore, the foreground-background segmentation algorithm has less precision with blur video.

7.6 Abnormally inactive during day

We conducted experiments to determine the sensitivity and specificity of raw and privacy-enhanced video of the low-activity-too-long detector. The amount of pixels that an MBR moves to be considered significant motion was trained by Actor 3 using only raw video in four situations, each 15 minutes long: person pace in main living area, person sweep in main, person lay still on couch, and person leave home. The detector was tested with Actor 1 and Actor 2 using raw and privacy-enhanced video in the same four situations. The average sensitivity and specificity for raw and each privacy-enhanced video were all 1.0.

7.7 Fall detection

We evaluated the sensitivity and specificity of our fall detection algorithm on 1 minute long privacy-enhanced video. We trained the fall detection algorithm with raw video then tested each privacy-enhanced video. The average sensitivity and specificity are shown in Table 2. Among the privacy-enhanced videos, the bounding-oval had the highest average sensitivity of 0.91 and the highest average specificity of 0.92, which are identical results to the raw video.

7.8 Comparison of privacy enhancements by goal performance

Table 3 contains the aggregated results of the experiments.

Table 2: Fall detection accuracy using MBR width trained on raw video and tested on each privacy enhancement. Higher is better.

Privacy enhancement	Average sensitivity	Average specificity
Raw	0.91	0.92
Blur	1.00	0.67
Silhouette	0.91	0.75
Bounding-oval	0.91	0.92
Bounding-box	0.82	0.92

Energy estimation and fall detection yielded the widest variation of results of the monitoring goals. For energy estimation, raw video had the highest accuracy and bounding-box had the highest fidelity. For fall detection, bounding-oval and raw video tied for the highest sensitivity and specificity. Of the remaining monitoring goals, the sensitivity and specificity were perfect 1.0, except blur video's sensitivity of 0.5 for in region too long. Using a privacy-enhanced video instead of raw video results in a 5% loss in energy estimation accuracy, but otherwise no major change in the ability to meet assistive monitoring goals.

We expected the goals to yield identical results for bounding-box and raw video, but raw video tended to outperform bounding-box video. Figure 4 shows a comparison between the background image, current frame, and foreground image for bounding-box video and raw video. The person in the recording entered the camera's view from the left, walked to the chair, then sat down. Figure 4(a) shows the background image for raw and bounding-box video. The background image of bounding-box contains noise caused by the solid box being learned into the background. This noise causes the current frame to differ less with the background image, which causes the foreground to be less accurate. The MBR around the foregrounds, shown in Figure 4(c), have different shapes and sizes.

Table 3: The aggregated results of the experiments for each assistive monitoring goal by privacy enhancement. Higher metrics are better. The privacy enhancements of silhouette, bounding-oval, and bounding-box performed as well as raw video.

Privacy enhancement	Energy estimation average accuracy	Energy estimation average fidelity	Room too long average sensitivity/ specificity	Arisen in morning sensitivity/ specificity	Region too long sensitivity	Region too long specificity	Abnormally inactive sensitivity/ specificity	Fall detection sensitivity	Fall detection specificity
Raw	90.9%	0.994	1.0 / 1.0	1.0 / 1.0	1.0	1.0	1.0 / 1.0	0.91	0.92
Blur	80.5%	0.991	1.0 / 1.0	1.0 / 1.0	0.5	1.0	1.0 / 1.0	1.00	0.67
Silhouette	85.0%	0.995	1.0 / 1.0	1.0 / 1.0	1.0	1.0	1.0 / 1.0	0.91	0.75
Bounding-oval	85.6%	0.994	1.0 / 1.0	1.0 / 1.0	1.0	1.0	1.0 / 1.0	0.91	0.92
Bounding-box	84.3%	0.997	1.0 / 1.0	1.0 / 1.0	1.0	1.0	1.0 / 1.0	0.82	0.92

8. Conclusion

We showed that eight common in-home assistive monitoring goals, including fall detection, energy estimation, and various basic scenario detections, could be achieved via automated video processing algorithms even in the presence of a variety of privacy enhancements. The privacy enhancements of a bounding-box, bounding-oval, or silhouette yielded nearly no loss in goal achievement. However, the common privacy enhancement of blurring did exhibit some loss in goal achievement, and thus its use in in-home assistive monitoring might be avoided, especially given the availability of the other privacy enhancements. Raw video had a fidelity of 0.994 for the goal of energy expenditure estimation, while silhouette had 0.995, bounding-oval had 0.994, and bounding-box had 0.997. The fall detection algorithm yielded the same sensitivity of 0.91 and specificity of 0.92 for raw and bounding-oval video, while silhouette had a sensitivity of 0.91 and specificity of 0.75, and bounding-box had a sensitivity of 0.82 and specificity of 0.92. The other 6 goals yielded perfect sensitivity and specificity for raw and privacy-enhanced video, with the exception of blur video's sensitivity of 0.5 in region too long.

The flexibility and decreasing costs of cameras may result in more in-home cameras. Future work involves dealing with real homes, multiple people, different camera angles, cameras and sensors working together, algorithms that adapt to the privacy enhancement, and more.

Figure 4: Comparison of raw video (left) and bounding-box privacy-enhanced video (right) for MBR determination: (a) The bounding-box video has more background noise. (b) The background noise causes the current frame to differ less with the background when a person is present. (c) The result is a larger, less accurate MBR for the bounding-box privacy-enhanced video.

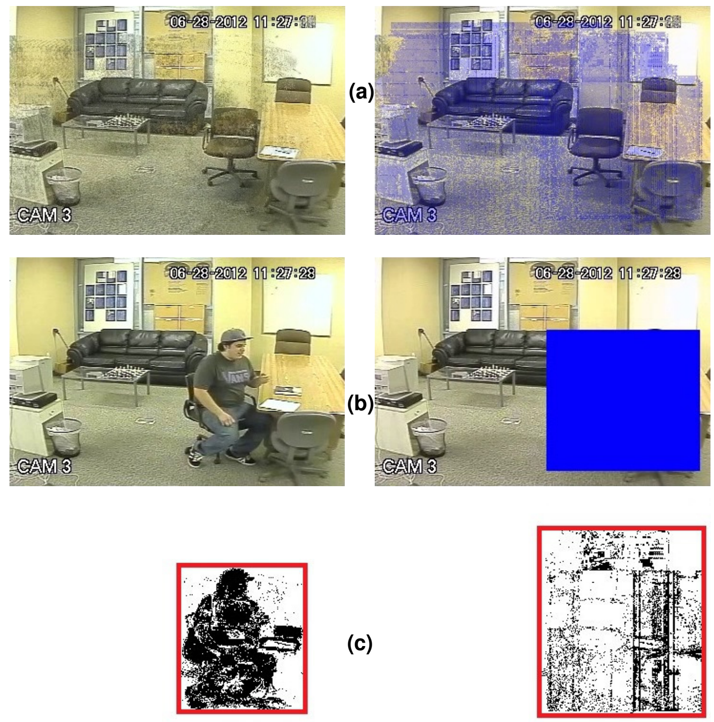

9. REFERENCES

[1] Anderson, D., 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, 2009.

[2] BeClose. http://beclose.com/. September 2012.

[3] Bourke, A.K., J.V. O'Brien, and G.M. Lyons. Evaluation of a threshold-based tri-axial accelerometer fall detection algorithm. Elsevier: Gate and Posture. Volume 26, pgs. 194-199, 2007.

[4] Bergmann, J.H.M. and A.H. McGregor. Body-Worn Sensor Design: What Do Patients and Clinicians Want? Annals of Biomedical Engineering. Volume 39, pgs. 2299-2312, 2011.

[5] Body Media Armbands. http://www.bodymedia.com/Shop/ Armband-Packages. August 2012.

[6] BodyBugg. http://www.bodybugg.com/. August 2012.

[7] Centers for Disease Control and Prevention. How much physical activity do older adults need? http://www.cdc.gov/physicalactivity/ everyone/guidelines/olderadults.html. August, 2012.

[8] Demiris, G., M.J. Rantz, M.A. Aud, K. D. Marek, H.W. Tyrer and M. Skubic, A.A. Hussam. Older adults' attitudes towards and perceptions of 'smart home' technologies: a pilot study. Medical Informatics and The Internet in Medicine, 2004.

[9] Edgcomb, A. and F. Vahid. Automated Fall Detection on Privacy-Enhanced Video. IEEE Engineering in Medicine and Biology Society, 2012.

[10] Edgcomb, A. and F. Vahid. Estimating Daily Energy Expenditure from Video for Assistive Monitoring. http://static.cs.ucr.edu/store/techreports/UCR-CS-2012-09260.pdf. November 2012.

[11] Edgcomb, A and F. Vahid. MNFL: The Monitoring and Notification Flow Language for Assistive Monitoring. ACM SIGHIT International Health Informatics Symposium (IHI), 2012.

[12] Edgcomb, A. and F. Vahid. Privacy perception and fall detection accuracy for in-home video assistive monitoring with privacy enhancements. ACM SIGHIT (Special Interest Group on Health Informatics) Record, 2012.

[13] Gator Tech Smart House - Smart Floor. http://www.icta.ufl. edu/pervasive-applications.htm#1. September 2012.

[14] Helal, S., W. Mann, H. El-Zabadani, J. King, Y. Kaddoura, and E. Jansen. The Gator Tech Smart House: A Programmable Pervasive Space. IEEE Computer. March 2005.

[15] KaewTraKulPong, P. and R. Bowden. An Improved Adaptive Background Mixture Model for Real-time Tracking with Shadow Detection. In Proc. 2nd European Workshop on Advanced Video Based Surveillance Systems, AVBS01, 2001.

[16] Kim, K. T.h. Chalidabhongse, D. Harwood and L. Davis. Real-Time Foreground-Background Segmentation using Codebook Model. Real-Time Imaging, Volume 11, Issue 3, June 2005, pgs. 172-185.

[17] Liu, L., M. Popescu, M. Skubic, T. Yardibi, and P. Cuddihy. Automatic Fall Detection Based on Doppler Radar Motion Signature. 5th International Conference on Pervasive Computing Technologies for Healthcare and Workshops, 2011.

[18] Motorola Homesight. http://www.smarthomeusa.com/Products/Homesight-USB-Kit/manuals/homesight-software-user-guide.pdf. September 2012.

[19] Nelson, M.E., W.J. Rejeski, S.N. Blair, P.W. Duncan, J.O. Judge, A.C. King, C.A. Macera, and C. Castaneda-Sceppa. Physical Activity and Public Health in Older Adults: Recommendation from the American College of Sports Medicine and the American Heart Association. Medicine and Science in Sports and exercise. Volume 39, pgs. 1435-1445, 2007.

[20] QuietCare. http://www.careinnovations.com/Products/QuietCare/. September 2012.

[21] Rashidi, P. and D.J. Cook. Keeping the Resident in the Loop: Adapting the Smart Home to the User. IEEE Transactions on Systems, Man and cybernetics, Part A: Systems and Humans, 2009.

[22] Rougier, C., J. Meunier, A. St-Arnaud, and J. Rousseau. Fall Detection from Human Shape and Motion History using Video Surveillance. 21st International Conference on Advanced Information networking and Applications Workshops, 2007.

[23] SmartHome. http://www.smarthome.com/. September 2012.

[24] Stauffer C, Grimson WEL. Adaptive background mixture models for real-time tracking. IEEE International Conference on Computer Vision and Pattern Recognition 1999; 2 : pg. 246–52.

[25] X10. http://www.x10.com. September 2012.

[26] Zhang, K., F.X. Pi-Sunyer, and C.N. Boozer. Improving Energy Expenditure Estimation for Physical Activity. Medicine and Science in Sports and Execise. Volume 36, pgs. 883-889, 2004.