

Non-Contact Millimeter-Wave Real-Time Detection and Tracking of Heart Rate on an Ambulatory Subject

Ilya V. Mikhelson, *Student Member, IEEE*, Philip Lee, *Student Member, IEEE*, Sasan Bakhtiari, *Senior Member, IEEE*, Thomas W. Elmer II, *Member, IEEE*, Aggelos K. Katsaggelos, *Fellow, IEEE*, and Alan V. Sahakian*, *Fellow, IEEE*

Abstract—This paper presents a solution to an aiming problem in the remote sensing of vital signs using an integration of two systems. The problem is that to collect meaningful data with a millimeter-wave sensor, the antenna must be pointed very precisely at the subject’s chest. Even small movements could make the data unreliable. To solve this problem, we attached a camera to the millimeter-wave antenna, and mounted this combined system on a pan/tilt base. Our algorithm initially finds a subject’s face and then tracks him/her through subsequent frames, while calculating the position of the subject’s chest. For each frame, the camera sends the location of the chest to the pan/tilt base, which rotates accordingly to make the antenna point at the subject’s chest. This paper presents a system for concurrent tracking and data acquisition with results from some sample scenarios.

Index Terms—Remote Sensing, Millimeter-Wave, Heart Rate, Patient Monitoring, Human Tracking.

I. INTRODUCTION

IN past research on remote subject monitoring, respiration and cardiac activity have been recorded remotely using different frequencies, such as 2.4 GHz [1], 10 GHz [2], 60 GHz [3], 94 GHz [4], [5], and even 228 GHz [6]. In each of these past studies, the subject was either seated or standing, but always stationary. Most often, the subject was also supported by either leaning against a wall when standing or the back of a chair when sitting to further reduce extraneous motion. Then, the antenna was manually aimed at the correct location on the chest to gather both respiratory and cardiac data.

However, there are numerous applications where this approach would not work. This kind of approach only works for short-term recordings and when the subject is deliberately staying still. For applications such as long-term patient monitoring and for security screening purposes such as for lie detection [7], the subject is typically not motionless. In those

situations, the ability to track the subject and aim at a particular location on his/her body becomes essential.

In patient monitoring, the current method to obtain physiological data is to attach a device to the body. This method can be uncomfortable for the patient, especially over long recordings. With the proposed system, the patient is free to move around uninhibited by instrumentation. Although there is other work in non-contact physiological data acquisition, it typically only works in a test environment, whereas the system reported herein could extend the technology to realistic situations.

Likewise, in a security setting such as covert detection of lying [7], the subject would not be standing still, as knowing that he/she was being monitored would defeat the purpose of the remote detection. Instead, the system would need to be able to follow the subject and still extract meaningful data despite the movement. Previous systems do not address this realistic scenario.

The novelty of our system lies in the fact that it automatically aims itself and follows the subject as he/she moves. The user is only required to turn on the system, and it finds a human subject and moves the entire antenna assembly to focus on the correct location on the subject’s chest. Furthermore, if the subject moves to a new location, the sensor will follow while maintaining aim at the proper location on the chest.

In most video-based person-tracking, the camera is stationary and only the subject is moving. This allows a broad range of techniques to be used [8], [9]. For our application, the camera moves with the antenna, and the motion of the image has to be accounted for in the algorithm. This requirement precluded the use of motion prediction algorithms such as Kalman filtering [10], since we would have to know the exact motion of the camera.

Also, in past research on person-tracking, it has been sufficient to find a rough location of the subject, denoted by an estimated bounding box in an image [9], [11], [12]. In this application, higher accuracy is necessary, as the antenna needs to be aimed precisely on the chest to obtain heart rate.

There has been past research on human tracking with a moving camera [13], [14]. However, aside from general tracking, we need an accurate estimate of the subject’s chest. Using a novel combination of the Viola-Jones algorithm [15] and CAM-Shift [16] tracking, we are able to estimate the position of the subject’s chest with high accuracy. In this way,

Manuscript received —; revised —. This work was supported in part by the U.S. Department of Energy.

I. V. Mikhelson, P. Lee, and A. K. Katsaggelos are with the Department of Electrical Engineering and Computer Science, Northwestern University, Evanston, IL 60208 USA (e-mail: i-mikhelson@u.northwestern.edu; philip-lee@u.northwestern.edu; aggk@eecs.northwestern.edu).

S. Bakhtiari and T. W. Elmer are with the Nuclear Engineering Division, Argonne National Laboratory, Argonne, IL 60439 USA (e-mail: bakhtiari@anl.gov; elmer@anl.gov).

*A. V. Sahakian is with the Department of Electrical Engineering and Computer Science and the Department of Biomedical Engineering, Northwestern University, Evanston, IL 60208 USA (e-mail: sahakian@eecs.northwestern.edu).



Fig. 1. Setup of system, including the antenna, the high-resolution camera, the depth camera, and the pan/tilt base.

data can be collected on a subject that is not just standing or sitting still.

This paper is organized as follows. In section II, we first present the millimeter-wave system and the detection and tracking algorithms. Then, we discuss how the separate systems are combined. Finally, in section III, we show our results for various test scenarios.

II. SYSTEM DETAILS

The system itself is a combination of several components: a 94 GHz vital sign monitoring sensor, a color camera, a depth camera, and a pan/tilt base (see Appendix for component details), which can be seen in Fig. 1. A computer coordinates the operation of the sensor assembly in order to produce meaningful results using a closed feedback loop. The color camera captures an image, detecting a human subject within. Based on the detected location of the chest, the entire assembly (see Fig. 1) moves to the appropriate position so as to point the antenna at the chest area. This movement is controlled by the pan/tilt control DAQ (see Appendix). As the antenna moves, the cameras constantly capture new images and process them to track the subject, with the depth camera aiding in accurate aiming by rejecting a lot of background noise and also by providing real-world coordinates as opposed to just pixel coordinates of the target. The details of this process are described in the subsequent sections.

A. Remote Detection of Cardiac Activity

Cardiac activity is processed by finding chest movements along the frontal axis corresponding to heartbeats. To obtain the chest displacement, a 94 GHz sensor was used [4], [17].

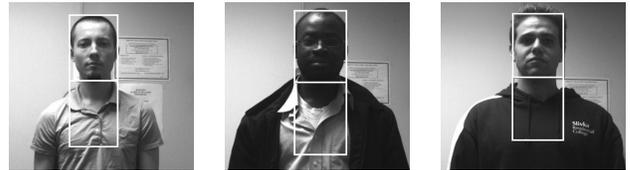


Fig. 2. Demonstration of proportions of head to chest.

With a 15.24 cm lens, the beam divergence was about 1.5° . This narrow beam means that precise aiming of the antenna is imperative to obtaining a good signal. The narrow beam also means that little extraneous data is gathered by the system. A meaningful signal is obtained when the antenna is aimed around the level of the nipples, and at the center of the chest. A very common estimate for humans, made by the Renaissance artist Leon Battista Alberti, was that the nipples are two head-lengths (length from top of head to chin) below the top of the head [18]. We checked this empirically on several subjects of different sizes, shown in Fig. 2.

To calculate a measure of displacement, an in-phase and a quadrature component of the reflected signal are calculated by mixing the original signal and a 90° phase-shifted version of the original signal with the reflected received signal to form an in-phase (I) and a quadrature (Q) component, respectively. Since the chest displacement d is modulated on the phase of the reflected signal, it can then be isolated using $d = (\lambda_0/4\pi) \arctan(Q/I)$, where λ_0 is the wavelength of the signal. The heartbeats can then be calculated using wavelet multiresolution decomposition and statistical techniques, as is detailed in [19].

To verify the locations of the estimated heartbeats, we also recorded an ECG concurrently with the displacement measurements. The details of the ECG system are in the Appendix. The subject had three electrodes attached to his body (left arm, right arm, right leg). The ECG served as a “gold standard” with which we verified our results.

B. Object Detection and Tracking

The vision aspect of our system consists of two main algorithms: detection and tracking. Detection refers to finding an instance of an object in the scene. Tracking refers to following a specific instance from frame to frame. The details and integration of the two follow.

1) Detection:

For detection we adopted the Viola-Jones detector [15], which is a particular classification system designed to be implemented in real time. We used Haar-like features [20], which are simply convolutions of binary-valued filters with rectangular support and allow for very efficient computation. Each feature can be calculated by evaluating the integral image [21] at a small number of points corresponding to the corners of the rectangular support regions of the filter.

Given a set of training images of fixed size (typically 20×20 pixels), Adaboost [22] is used as the learning algorithm to combine a large set of single-feature classifiers into a strong classifier. Since the number of weak classifiers included in the strong classifier may be large, it is computationally

inefficient to classify each possible subwindow in a new image. Therefore, the weak classifiers are grouped into stages such that the first stage rejects the majority of subwindows in the image. Processing continues until a stage rejects the subwindow as not containing the object, or until the final stage is completed, and the subwindow is considered to contain the object. In this way, minimal processing is consumed for most subwindows, and results in a real time algorithm for object detection.

It is important to notice that even though the training images are fixed to a specific size, the algorithm can handle multiple scales simply by scaling the input image to different sizes. Then, the Viola-Jones detector can be applied to these images to determine the location and approximate scale of the object.

2) Tracking:

For tracking, we use an implementation of the Continuously Adaptive Mean-Shift (CAM-Shift) algorithm [16], which is an extension of the Mean-Shift algorithm [23]. This method is one of the few acceptable methods for our implementation due to the extra complexity added by a moving camera. Mean-Shift operates as follows. Given a way to generate a probability map of where an object is located in the image and an initial guess, the algorithm iterates and changes the guess so as to locally maximize the probability distribution, called a *backprojection*. The backprojection is calculated here from the color histogram of the object. Given the object's histogram, for each small window another histogram is generated and compared to that of the object. The greater the similarity, the greater the value of the backprojection at that location. CAM-Shift simply extends this to include not only location, but also scale and rotation, and performs the iterations in this 4-dimensional space. In fact, the reason we chose CAM-Shift over the less computationally intensive Mean-Shift is that, for this application, the scale information is crucial, and the angle information is very useful as well.

We further extended the utility of the CAM-Shift algorithm by implementing the ABC-Shift algorithm [24]. Instead of having just one tracking window, we also included an outer window twice the size of the tracking window. If an object with a similar color distribution enters the outer window, those colors are de-emphasized in the resulting backprojection. Therefore, the tracking window tends to drift less onto the similarly colored object, even if it gets very close to the tracked object.

3) Depth Data:

Using the Microsoft Kinect™, we were also able to gather depth data. We first performed a calibration between the high-resolution camera and the depth camera using a custom GUI created with the OpenCV library [25]. Then, we used the depth data to segment the image in depth. Only pixels that are within 0.5 meters of the subject of interest are used for processing. This eliminated much of the noise in the image and also helped with identity preservation, as only similar objects within 0.5 meters of the subject could potentially interfere.

In addition to depth segmentation, the depth data made it possible for us to know real-world coordinates as opposed to pixel coordinates. This allowed us to perform velocity estimation and thus motor motion compensation in our pro-

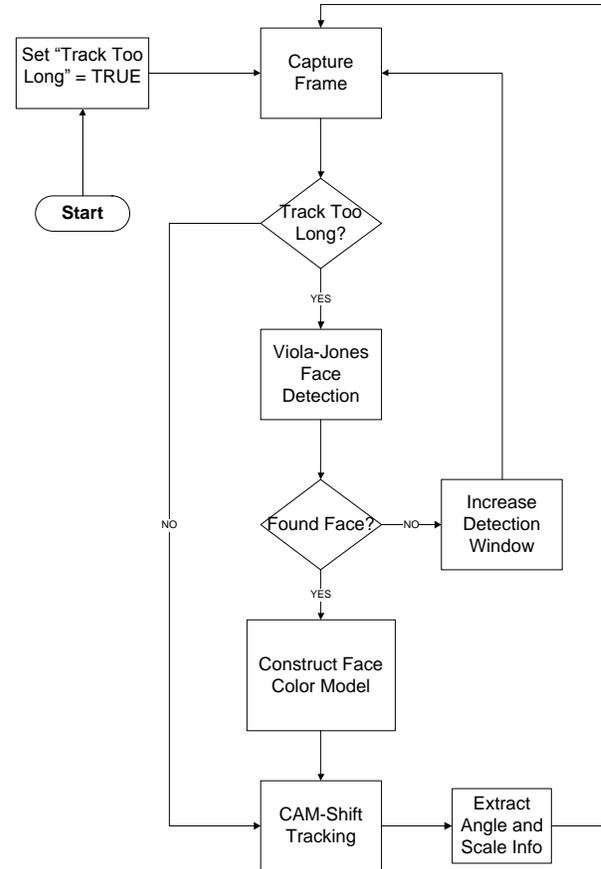


Fig. 3. Flow Chart of detection and tracking algorithms.

cessing. This means that we were able to account for the movement between consecutive frames and still accurately aim the sensors. In addition, we were able to aim the antenna very precisely, since we know its exact translation from the high-resolution camera in real-world coordinates.

4) Combination of Detection and Tracking:

To find and follow a person from frame to frame, we use a combination of the detection and tracking algorithms described above and illustrated in Fig. 3. We wrote the implementation using the OpenCV library [25] in C++.

For detection, we used two cascades of classifiers: one that recognizes a frontal head/shoulders combination, and one that recognizes the profile of a face. This combination is robust enough to detect a person in most reasonable poses.

Given a cascade of classifiers for a particular object, the detector returns several hypotheses for location and scale of instances of the object. To initialize the algorithm, we run both cascades in the detector, preferring head/shoulders detections, and pick the strongest hypothesis as the target object. For subsequent detections, we make the assumption that the object cannot instantly jump from one location to another, and therefore bound the search window for future detections to a small radius around the last detection, which is initially twice the width of the bounding box of the face. As we accumulate detection failures in this search window, we expand it until the search window is the entire image or we have another successful detection, in which case the search

radius is reset. This heuristic helps to maintain the identity of the tracked object, meaning that even if multiple instances are present, it will tend to detect the desired object only.

At this point, we have the location and scale of the face of the person we wish to track. In order to do tracking, we must first provide a feature model of the face. We use a model based on a hue histogram within the initial detection window because hue is invariant to lighting conditions, which are very likely to change from frame to frame. However, the calculation of the hue is stable only when a pixel is not too dark or bright. For this reason, we only include those pixels that have a saturation of at least 100/256 and whose value is between 65/256 and 200/256. These values were chosen empirically from our own experiments to provide the best results.

The hue histogram model is used to compute the backprojection at each frame needed by the CAM-Shift tracker. However, in order for the CAM-Shift algorithm to be stable, it was necessary to first apply a small-radius (3 pixels) median filter to the backprojection to get rid of the salt and pepper noise, which had created tiny local maxima for the CAM-Shift iteration, causing scale, angle, and location drift. Since the backprojection tends to be piecewise constant, we applied a 7x7 Gaussian filter ($\sigma = 1.55$) to further reduce drifting by helping CAM-Shift to converge.

After creating this backprojection, CAM-Shift tracking occurs exclusively for a certain number of frames, after which we apply the detector again in an attempt to correct any further drift that may have occurred. Initially, the time between detections is 2 frames. If detection fails (*i.e.*, there is no clear frontal or profile face in the frame) 3 consecutive times, then CAM-Shift runs for an increased amount of time. The time between detection attempts increases each time the detector fails. This approach helps to prevent the CPU-intensive detector from running too often when it is unable to make a good hypothesis. However, even if there are very few detections, we do not let CAM-Shift run for more than 40 frames so that CAM-Shift does not become the dominant mode of operation in the algorithm.

C. Integration of Vision and Cardiac Pattern Detection

First, we compensate for the offset between the optical centers of the camera and the antenna by using real-world coordinates provided by the Kinect and the known offset between the devices.

Using the accurate location of the face provided by the detection and tracking algorithms, we can find the location of the subject's chest heuristically. Using the scale of the subject's face estimated by the algorithm, the location of the chest can be estimated based on the fact that human body parts tend to have certain proportional relationships. We therefore used the heuristic described in section II-A with a slight modification to account for realistic situations, and aimed the antenna 1.5 head-lengths below the top of the head along the vertical midline of the bounding box of the head. We did this because the bounding box includes the subject's neck, which makes the aim point too low. Therefore, we aim higher than twice the head length below the top of the head. If the

TABLE I
DESCRIPTION OF SUBJECTS.

Subject #	Gender	Age	Height (cm)	BMI
1	F	20	164	18.3
2	M	24	180	26.5
3	M	25	173	22.8
4	M	26	168	20.2
5	M	57	173	25.5

subject is wearing clothing covering his/her neck, the aim point is higher than at nipple level, but that did not cause a problem in data acquisition. When we tested this with the actual antenna, it captured meaningful data (*i.e.*, the estimated heart rate was close to the actual heart rate calculated from the ECG). Therefore, this was the estimate that we used when aiming the antenna.

The angle estimated by the algorithm is also of use when the subject's head is not vertically oriented. If the tilt of the head is greater than 45°, the aim point is at one length of the semi-major axis of the bounding ellipse below the bottom (near the subject's neck) of the ellipse, as shown in Fig. 13.b.

It should be noted that our algorithm does indeed calculate the size of the entire head, even though only face detection is being performed. This is due to the fact that we track based on a hue histogram, which includes all skin-colored pixels and not just the face. However, if the subject has hair covering his/her forehead, the system still aims the antenna at an appropriate location on the chest, although it may be a little higher due to an underestimation of the size of the head.

With this implementation, we are able to find and track a person while recording his/her cardiac patterns automatically. Once the system is turned on, it locates a person, aims the antenna at his/her chest, and follows the subject as long as he/she stays within the camera's view.

To avoid jerking of the system caused by small movements of the subject, the pan/tilt base only moves when the estimated location of the chest has moved more than ten pixels from the previous location. This keeps the antenna aimed at a reasonable location on the chest while preventing abrupt movements of the system that could make the data unusable.

There are a couple of practical limitations to this system. First, the subject cannot move too quickly, because the pan/tilt head's speed for the existing system is limited. Another limitation is due to the fact that we use a hue histogram to model the target. This means that if a similarly colored object occludes the subject's face, the CAM-Shift algorithm may drift on to this new object until detection can correct the error.

III. RESULTS AND DISCUSSION

In the following demonstrations, we show how the system performs in several scenarios with movement. The main measure of accuracy is whether or not we obtain meaningful data (as defined in section II-C) with the millimeter-wave interferometer. The subjects in the following scenarios are presented in table I. Testing was done with approval from Northwestern University's Office for the Protection of Research Subjects (IRB Project Number: STU00051704).

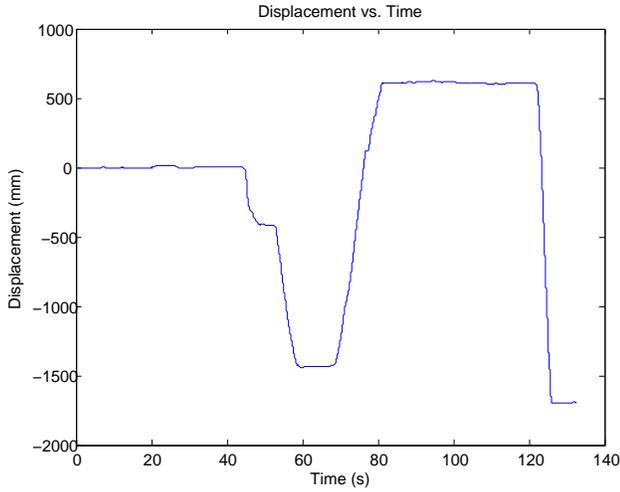


Fig. 4. Displacement waveform for subject #4.

A. Cardiac Pattern Detection on Ambulatory Subjects

Each subject starts off seated against a chair back at 3.5 m from the sensor. While seated, the subjects holds his/her breath for about 30 seconds and then breathes regularly for about 10 seconds. Then, the subject stands up and walks slowly toward the sensor up to a distance of 2 m from the sensor. The subject stands there for about 10 seconds and subsequently moves backward to 4 meters from the sensor. There, the subject again holds his/her breath for about 30 seconds and breathes regularly for about 10 seconds. Finally, the subject moves toward the sensor more quickly until he/she is 2 m from the sensor. A sample plot of this movement can be seen in Fig. 4 (subject #4).

Each subject, along with a clip of data, can be seen in Figs. 5, 6, 7, 8, and 9. On the left side is the subject. The thick white circle indicates the aim location of the antenna and the gray bounding box surrounding the face was calculated as the circumscribed rectangle of the CAM-Shift ellipse (white). On the top right is a plot of a section of the displacement signal for the subject with calculated heartbeats indicated by stems. On the bottom right is the ECG that was recorded concurrently with data collection, where the calculated heartbeat locations are indicated as circles and QRS complex locations are shown by crosses.

Individual heartbeats can clearly be seen when there is no movement due to breathing or body movement (Figs. 5, 6). It can be seen that the estimated heartbeats [19], shown by stems, have a one-to-one correspondence with the subject’s ECG, where the estimated heartbeats are shown as circles above the ECG and the crosses indicate QRS complex locations. This indicates that the antenna is aimed correctly and that the data gathered is meaningful.

Using these five data sets, we performed a leave-one-out cross-validation, where we created a precision-recall plot by varying thresholds in the algorithm [19] for 4 subjects and used the thresholds that created the point farthest from the origin to test the remaining subject. As can be seen in Fig.

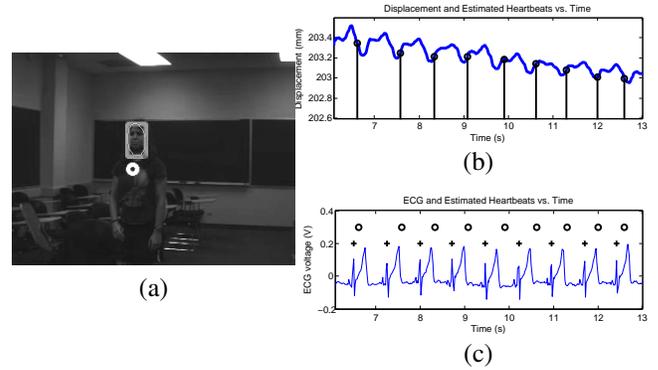


Fig. 5. Subject #1: a) System locked on to subject’s chest, b) Calculated displacement of location indicated by thick white circle in (a), with calculated heartbeats indicated by stems, and c) Recorded ECG with calculated heartbeat locations shown as circles and QRS complex locations shown by crosses.

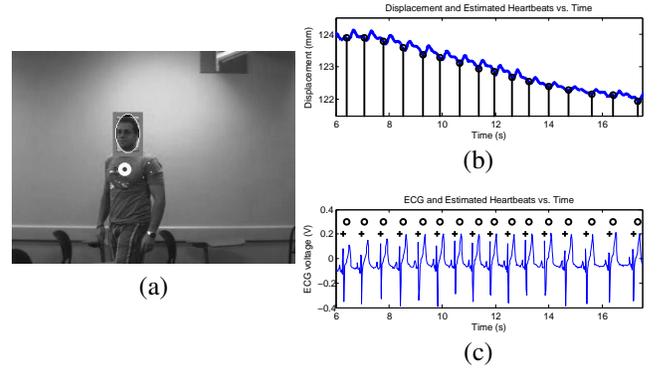


Fig. 6. Subject #2: a) System locked on to subject’s chest, b) Calculated displacement of location indicated by thick white circle in (a), with calculated heartbeats indicated by stems, and c) Recorded ECG with calculated heartbeat locations shown as circles and QRS complex locations shown by crosses.

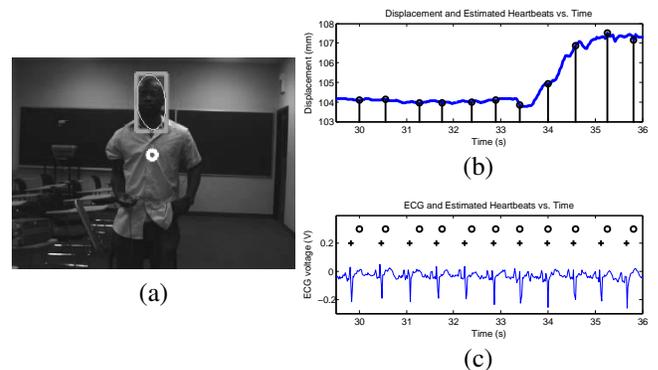


Fig. 7. Subject #3: a) System locked on to subject’s chest, b) Calculated displacement of location indicated by thick white circle in (a), with calculated heartbeats indicated by stems, and c) Recorded ECG with calculated heartbeat locations shown as circles and QRS complex locations shown by crosses.

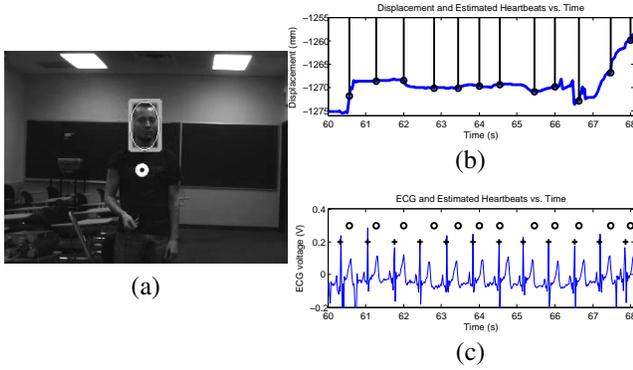


Fig. 8. Subject #4: a) System locked on to subject's chest, b) Calculated displacement of location indicated by thick white circle in (a), with calculated heartbeats indicated by stems, and c) Recorded ECG with calculated heartbeat locations shown as circles and QRS complex locations shown by crosses.

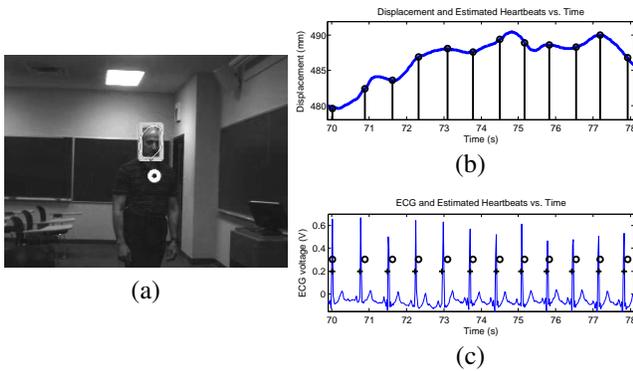


Fig. 9. Subject #5: a) System locked on to subject's chest, b) Calculated displacement of location indicated by thick white circle in (a), with calculated heartbeats indicated by stems, and c) Recorded ECG with calculated heartbeat locations shown as circles and QRS complex locations shown by crosses.

TABLE II
RESULTS FOR DATA WITH MOVEMENT.

Subject #	TP	FP	FN	Recall	Precision
1	108	18	5	95.6%	85.7%
2	134	16	5	96.4%	89.3%
3	275	4	53	83.8%	98.6%
4	176	24	4	97.8%	88.0%
5	145	16	10	93.6%	90.1%
Total	838	78	77	91.6%	91.5%

10 (a precision-recall curve for subjects #2, 3, 4, and 5), the curve passes close to the top right corner (where the top right corner is perfect detection). The results of the cross-validation are given in table II.

As can be seen, subject #3 had a much larger number of false negatives than false positives. This is because the calculated heartbeats were not placed close enough to the "true" heartbeat locations because they were interpolated over a large area of "unconfident" heartbeats. We are therefore continuing to work on ways to guarantee that there are no long periods without any "confident" heartbeats.

B. Robustness of Detection and Tracking

We also demonstrate that depth motion of the subject in relation to the camera (*i.e.*, movements toward and away from

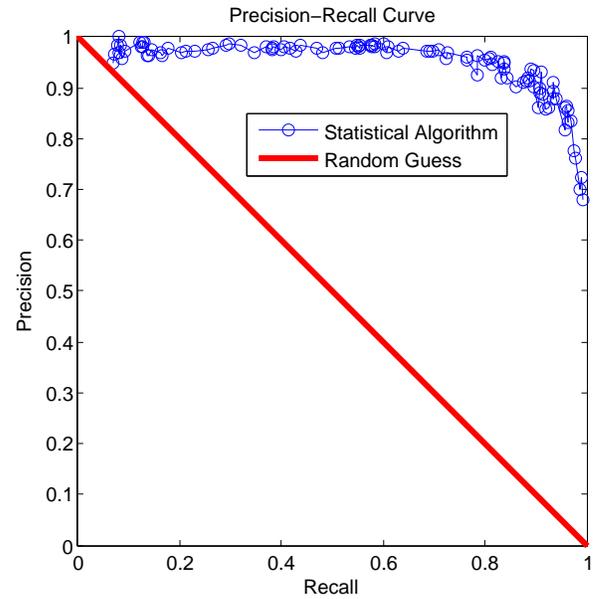


Fig. 10. Precision-Recall curve for movement data.



Fig. 11. a) System locked on to chest when subject is initially far away, and b) System locked on to chest when subject is moving closer.

the camera) can be compensated for with our algorithm. Since we gather scale information with our algorithm, we are able to heuristically determine the location of the chest regardless of depth. Therefore, the subject can move laterally and axially with respect to the antenna and the system tracks him/her accurately. Fig. 11.a shows the subject 9 meters away from the system. The system is locked on and gathering data. In Fig. 11.b, the subject has moved to 2.5 meters away from the system. The antenna is still pointing accurately at his chest. It should be noted that the antenna is aiming at the correct location regardless of scale. This is important because even though the size of the subject in the image has changed, the system is able to aim accurately.

Another important aspect of our system is that it tends to maintain identity and is robust to partial occlusion. Fig. 12 shows another person approaching the subject of interest and the system stays focused on the correct person. This is true even when the other person gets very close. This is because the detection window stays close to the first subject's face, and when the other person gets close, the ABC-Shift algorithm de-emphasizes similar colors so as to minimize any drift. However, if the other person gets inside the search window, there can be an error if the detection algorithm deems the new



Fig. 12. System maintaining identity when another human is close to the subject of interest.



Fig. 13. a) System handling occlusion of area that is being tracked, and b) System handling a pose change while maintaining proper aim at the subject's chest.

face as a “better” face (*i.e.*, better matching the features used to find a face) and if the face is similar in color. In that case, the system may shift to following the new subject.

In the case of partial occlusion, the system mitigates potential errors well using the ABC-Shift algorithm. Fig. 13.a shows the subject's face partially occluded by his hand holding a beverage can. The system, however, stays focused on the subject and the antenna points accurately at his chest without his hand causing tracking drift.

The system can also handle common pose changes and still aim accurately at the subject's chest. Specifically, if the subject moves his head to any angle other than upright, the system is still able to heuristically aim at his chest. These movements are common, such as when looking at a watch, looking up, or tilting the head for any reason. An example of this can be seen in Fig. 13.b, where the subject has leaned slightly to pick up an object.

IV. CONCLUSION

The integration of a millimeter-wave sensor, two cameras, and a pan/tilt base has been successfully demonstrated. A subject was detected and tracked by the system, while his/her chest displacement data was collected remotely with a millimeter-wave sensor. From this data, locations of the heart beats could be extracted. Future work in this area will include extraction of the respiratory rate from data when the subject is breathing normally.

The system was able to handle both lateral and axial motion of the subject while maintaining a good aim at the chest. We were able to process the data collected from subjects in motion and to accurately detect their heartbeats. A future goal is to make the feedback loop better by using the system motion information, such as rotational velocity, in the computer

processing to make the data gathered while the subject is in motion cleaner.

There are some situations in which the system fails, though. Since the tracking is based on a color model of the subject, other similarly colored objects may interfere with the tracking. However, detection along with ABC-Shift, as well as depth segmentation, most often fixes any problems before they get out of control. To combat potentially problematic scenarios, we intend to extend the algorithm to include more discriminating models, such as eigenfaces [26] and face template matching, to further reduce tracking drift.

APPENDIX SYSTEM COMPONENT DETAILS

Millimeter-wave Sensor

- 94 GHz Cavity tuned Gunn diode oscillator and quadrature mixer/detector
- Continuous-wave mode of operation
- Gaussian quasi-optical antenna with a 15.24 cm lens

Color Camera

- Sony
- Model: XCD-X710CR
- Resolution: 1024x768 RGB @ 30fps

Depth Camera

- Microsoft Kinect
- Resolution: 640x480 XYZ

Pan/tilt base

- Model: IPTU1
- Pan speed: 6 degrees/second
- Tilt speed: 3 degrees/second

Pan/tilt Control DAQ

- National Instruments
- USB-6212 16-bit M Series MIO DAQ

Data Acquisition DAQ

- National Instruments
- USB-9239 24-bit analog input module

ECG Monitor

- Tektronix 408 ECG Monitor

ACKNOWLEDGMENTS

The authors would like to thank Cindy Solomon and Oyinlolu Adeyanju for providing data to strengthen the results.

REFERENCES

- [1] D.R. Morgan and M.G. Zierdt, “Novel signal processing techniques for Doppler radar cardiopulmonary sensing,” *Signal Processing*, vol. 89, no. 1, pp. 45–66, 2009.
- [2] O.B. Lubecke, P.W. Ong, and VM Lubecke, “10 GHz Doppler radar sensing of respiration and heart movement,” in *Bioengineering Conference, 2002. Proceedings of the IEEE 28th Annual Northeast*. IEEE, 2002, pp. 55–56.
- [3] D. Obeid, S. Sadek, G. Zaharia, and G. El-Zein, “Non-contact heartbeat detection at 2.4, 5.8 and 60 GHz: A comparative study,” *Microwave Opt. Technol. Lett.*, vol. 51, no. 3, pp. 666–669, 2009.
- [4] I. V. Mikhelson, S. Bakhtiari, T. W. Elmer, II, and A. V. Sahakian, “Remote sensing of heart rate and patterns of respiration on a stationary subject using 94-ghz millimeter-wave interferometry,” *Biomedical Engineering, IEEE Transactions on*, vol. 58, no. 6, pp. 1671–1677, June 2011.

- [5] S. Bakhtiari, S. Liao, T. W. Elmer, N. Gopalsami, and A.C. Raptis, "A Real-time Heart Rate Analysis for a Remote Millimeter Wave I-Q Sensor," *Accepted to IEEE Transactions on Biomedical Engineering*.
- [6] D.T. Petkie, C. Benton, and E. Bryan, "Millimeter wave radar for remote measurement of vital signs," in *Radar Conference, 2009 IEEE*. IEEE, 2009, pp. 1–3.
- [7] J. Geisheimer and EF Grenaker III, "A non-contact lie detector using radar vital signs monitor (RVSM) technology," *Aerospace and Electronic Systems Magazine, IEEE*, vol. 16, no. 8, pp. 10–14, 2002.
- [8] A. Elgammal, D. Harwood, and L. Davis, "Non-parametric model for background subtraction," *Computer VisionECCV 2000*, pp. 751–767, 2000.
- [9] R. Bodor, B. Jackson, and N. Papanikolopoulos, "Vision-based human tracking and activity recognition," in *Proc. of the 11th Mediterranean Conf. on Control and Automation*. Citeseer, 2003, vol. 1.
- [10] R.E. Kalman, "A new approach to linear filtering and prediction problems," *Journal of basic Engineering*, vol. 82, no. 1, pp. 35–45, 1960.
- [11] B. Wu and R. Nevatia, "Tracking of multiple, partially occluded humans based on static body part detection," in *Computer Vision and Pattern Recognition, 2006 IEEE Computer Society Conference on*. IEEE, 2006, vol. 1, pp. 951–958.
- [12] D. Roth, P. Doubek, and L.V. Gool, "Bayesian pixel classification for human tracking," in *IEEE Workshop on Motion and Video Computing, 2005. WACV/MOTIONS'05 Volume 2, 2005*, pp. 78–83.
- [13] Yiming Ye, J.K. Tsotsos, K. Bennet, and E. Harley, "Tracking a person with pre-recorded image database and a pan, tilt, and zoom camera," in *Visual Surveillance, 1998. Proceedings., 1998 IEEE Workshop on*, Jan. 1998, pp. 10–17.
- [14] A. Mian, "Realtime face detection and tracking using a single pan, tilt, zoom camera," in *Image and Vision Computing New Zealand, 2008. IVCNZ 2008. 23rd International Conference, 2008*, pp. 1–6.
- [15] P. Viola and M. Jones, "Rapid object detection using a boosted cascade of simple features," in *IEEE Computer Society Conference on Computer Vision and Pattern Recognition*. Citeseer, 2001, vol. 1.
- [16] G.R. Bradski, "Computer vision face tracking for use in a perceptual user interface," 1998.
- [17] S. Bakhtiari, T.W. Elmer, N.M. Cox, N. Gopalsami, A.C. Raptis, S. Liao, I. Mikhelson, and A.V. Sahakian, "Compact millimeter-wave sensor for remote monitoring of vital signs," *Instrumentation and Measurement, IEEE Transactions on*, vol. 61, no. 3, pp. 830–841, march 2012.
- [18] J.A. Aiken, "Leon Battista Alberti's System of Human Proportions," *Journal of the Warburg and Courtauld Institutes*, vol. 43, pp. 68–96, 1980.
- [19] I. V. Mikhelson, S. Bakhtiari, T. W. Elmer, II, and A. V. Sahakian, "Remote sensing of patterns of cardiac activity using statistical methods and non-linear optimization," *submitted*.
- [20] R. Lienhart and J. Maydt, "An extended set of haar-like features for rapid object detection," in *IEEE ICIP*. Citeseer, 2002, vol. 1, pp. 900–903.
- [21] F.C. Crow, "Summed-area tables for texture mapping," in *Proceedings of the 11th annual conference on Computer graphics and interactive techniques*. ACM, 1984, pp. 207–212.
- [22] Y. Freund and R. Schapire, "A decision-theoretic generalization of on-line learning and an application to boosting," in *Computational learning theory*. Springer, 1995, pp. 23–37.
- [23] D. Comaniciu and P. Meer, "Mean shift: A robust approach toward feature space analysis," *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 24, no. 5, pp. 603–619, 2002.
- [24] R. Stolkin, I. Florescu, and G. Kamberov, "An adaptive background model for camshift tracking with a moving camera," in *Proc. International Conference on Advances in Pattern Recognition, 2007*, pp. 147–151.
- [25] G. Bradski and A. Kaehler, *Learning OpenCV: Computer vision with the OpenCV library*, O'Reilly Media, 2008.
- [26] M. Turk and A. Pentland, "Eigenfaces for recognition," *Journal of cognitive neuroscience*, vol. 3, no. 1, pp. 71–86, 1991.



Ilya V. Mikhelson (S'08) received the B.S. and M.S. degrees in electrical engineering from Northwestern University, Evanston, IL, in 2009 and 2011.

He is currently a Ph.D. candidate in electrical engineering at Northwestern University. His research interests include digital signal processing, computer vision, and remote patient monitoring.



Philip Lee (S'11) received the B.S. degree in mathematical sciences from Clemson University, Clemson, SC, in 2008, and the M.S. degree in electrical engineering & computer science from Northwestern University, Evanston, IL, in 2011.

He is currently a Ph.D. candidate in electrical engineering & computer science at Northwestern University, Evanston, IL. His research interests include computer vision, pattern recognition, and internet-scale computing.



Sasan Bakhtiari (M'91-SM'94) received the B.S.E.E. degree from the Illinois Institute of Technology, Chicago, in 1983, the M.S.E.E. from the University of Kansas, Lawrence, in 1987, and the Ph.D. degree in electrical engineering from Colorado State University, Fort Collins, in 1992.

From 1984 to 1987, he was with the Radar Systems and Remote Sensing Laboratory, University of Kansas. From 1998 to 1992, he was with the Microwave Nondestructive Testing Laboratory, Colorado State University. In 1993, he joined Argonne

National Laboratory, Argonne, IL, where he is currently a Principal Electrical Engineer with the System Technologies and Diagnostics Department at the Nuclear Engineering Division. He is also the Manager of the Nondestructive Evaluation (NDE) Section. He has authored or coauthored numerous technical publications in the area of sensors and NDE.



Thomas W. Elmer II (M'07) received the B.S. degree in physics (with minors in math and computer science) from La Sierra University, Riverside, CA, in 1998, and the M.S. degree from the University of Illinois at Chicago, in 2004.

While with La Sierra University, he was with the Physics Department, where he was involved with writing and maintaining programs to run laboratory experiments. He has also lectured on astronomy and gravitational physics for the Physics Department.

His senior research project was with the Health Physics Department, Loma Linda University Medical Center, for which he analyzed the activation of Cerrobend(R) metal by 250-MeV protons at the hospitals Proton Facility. In 1999, he joined Argonne National Laboratory, Argonne, IL, as a Student Intern, eventually remaining as a Software Engineering Associate for the System Technologies and Diagnostics Department, Nuclear Engineering Division. He writes programming for modeling, motion control, data acquisition, and data analysis in the microwave, millimeter-wave, and terahertz sensors laboratories.

Mr. Elmer was the recipient of the 2007 Research and Development (R&D) 100 Award presented by R&D Magazine and currently has two U.S. patents.



Aggelos K. Katsaggelos (S'80-M'85-SM'92-F'98) received the Diploma degree in electrical and mechanical engineering from the Aristotelian University of Thessaloniki, Greece, in 1979, and the M.S. and Ph.D. degrees in Electrical Engineering from the Georgia Institute of Technology, in 1981 and 1985, respectively.

In 1985, he joined the Department of Electrical Engineering and Computer Science at Northwestern University, where he is currently a Professor holder of the AT&T chair. He was previously the holder of the Ameritech Chair of Information Technology (1997- 2003). He is also the Director of the Motorola Center for Seamless Communications, a member of the Academic Staff, NorthShore University Health System, an affiliated faculty at the Department of Linguistics and he has an appointment with the Argonne National Laboratory.

He has published extensively in the areas of multimedia signal processing and communications (over 180 journal papers, 400 conference papers and 40 book chapters) and he is the holder of 19 international patents. He is the co-author of Rate- Distortion Based Video Compression (Kluwer, 1997), Super-Resolution for Images and Video (Claypool, 2007) and Joint Source-Channel Video Transmission (Claypool, 2007).

Among his many professional activities Prof. Katsaggelos was Editor-in-Chief of the IEEE Signal Processing Magazine (1997-2002), a BOG Member of the IEEE Signal Processing Society (1999-2001), and a member of the Publication Board of the IEEE Proceedings (2003-2007). He is a Fellow of the IEEE (1998) and SPIE (2009) and the recipient of the IEEE Third Millennium Medal (2000), the IEEE Signal Processing Society Meritorious Service Award (2001), the IEEE Signal Processing Society Technical Achievement Award (2010), an IEEE Signal Processing Society Best Paper Award (2001), an IEEE ICME Paper Award (2006), an IEEE ICIP Paper Award (2007) and an ISPA Paper Award (2009). He was a Distinguished Lecturer of the IEEE Signal Processing Society (2007-2008).



Alan V. Sahakian (S'84-M'84-SM'94-F'07) received the Ph.D. degree in electrical engineering from the University of Wisconsin, Madison, in 1984.

In 1984, he joined Northwestern University, Evanston, IL, where he is currently a Professor and the Chair in the Department of Electrical Engineering and Computer Science, and a Professor in the Department of Biomedical Engineering, and is also a Member of the Academic Affiliate Staff of North Shore University Health System (Evanston Hospital, Evanston, IL). He is the author or coauthor of approximately 200 papers, abstracts, patents, and book sections. His current research interests include the areas of cardiac electrophysiology, RF, microwave and millimeter-wave imaging and remote sensing and patient monitoring methods, and irreversible electroporation methods for tumor ablation.

Dr. Sahakian has served as the Vice-President for Publications and Technical Activities for the Engineering in Medicine and Biology Society. He is a Fellow of the American Institute for Medical and Biological Engineering.