Collecting Heart Rate Using a High Precision, Non-Contact, Single-Point Infrared Temperature Sensor

Laura Boccanfuso, Eva Juarez Perez, Myra Robinson, and Jason M. O'Kane

University of South Carolina, South Carolina SC 29208, USA,
boccanfu@cse.sc.edu

Abstract. Remotely detecting the physiological state of humans is becoming increasingly important for rehabilitative robotics (RR) and socially assistive robotics (SAR) because it makes robots better-suited to work more closely and more cooperatively with humans. This research delivers a new non-contact technique for detecting heart rate in real time using a high precision, single-point infrared sensor. The proposed approach is an important potential improvement over existing methods because it collects heart rate information unencumbered by biofeedback sensors, complex computational processing or high cost equipment. We use a thermal infrared sensor to capture subtle changes in the sub-nasal skin surface temperature to monitor cardiac pulse. This study extends our previous research in which breathing rate is automatically extracted using the same hardware. Experiments conducted to test the proposed system accuracy show that in 72.7% of typical cases heart rate was successfully detected within 0-9 beats per minute as measured by root-mean-square error.

1 INTRODUCTION

The growing use of robots in rehabilitative therapy and in socially assistive applications has brought to focus the need to make human-robot interactions as natural and beneficial to the humans using them as possible. A multitude of human-robot applications stand to greatly benefit from a small, inexpensive system capable of delivering accurate heart rate data remotely. For example, existing rehabilitative robotic (RR) systems such as shoulder, wrist, hand and ankle robots which exercise a targeted muscle or muscle group [1] may sense a user's level of strength to determine how much assistance to provide but would be further improved by having a non-contact method for obtaining their user's heart rate and overall stress level during therapy. Home-use robots working in close proximity to the elderly would be additionally beneficial to their users if they were able to remotely collect and transmit critical vital signs indicative of their physical condition to a family member or service provider. Further, robots used in socially assistive robotics (SAR) that are capable of monitoring the heart rate of children with social or developmental disabilities (especially those who are non-verbal) would be better-suited for use in therapy since their behavior could be adapted based on the perceived stress-state of the child. Should the child become stressed during the course of the therapy, he or she may not be able to immediately communicate this fact to the therapist or teacher. A robot that is continually collecting information about the child's heart rate can detect subtle shifts in his or her emotional state and alert the child's therapist before the child's frustration escalates. Detecting heart rate remotely is a necessary next step towards fully realizing this potential.

A variety of methods have been used to collect data about a user's emotional or stress state including measuring the amount of eye contact, body pose, number, quality and content of verbal utterances, and several physiological indicators such as galvanic skin response, electroencephalography (EEG), breathing and heart rate. Galvanic skin response measures changes in the electrical conductance of skin [2] while EEG is used to measure the voltage fluctuations resulting from ionic current flows within the neurons of the brain [3]. Capturing heart rate has traditionally employed contact modalities for obtaining data, although non-contact methods have more recently been explored.

Contact approaches include the use of electrocardiograms (ECG) which require the user to be fitted with sensors and a variety of recently introduced cardiograph applications for portable electronic devices which require the user to place their finger over a small, onboard camera to detect subtle changes in skin color. These devices typically deliver accurate heart rate data, but are generally not suitable for mobile applications, where lighting conditions are not consistent, for people who are averse to wearing sensors, and when the use of contact sensors is otherwise impractical. In addition, although solutions exist using non-contact methods such as radar and doppler modalities, these approaches rely on high-cost equipment and collecting and analyzing very large amounts of data at a high processing cost.

This research presents a new non-contact heart rate measurement technique suitable for most RR and SAR applications. Changes in the sub-nasal skin surface temperature are tracked and a heart rate in beats per minute (bpm) is automatically calculated. This study extends our previous research [4] in which breathing rate is computed using the same hardware we use for this study. Key improvements to our software have yielded a four-fold increase in the number of samples collected per second and the implementation of a Discrete Wavelet Transform (DWT) for automatically calculating the heart rate. The novel contribution of this paper is a simple, robust, low cost approach for remotely collecting and monitoring heart rate.

The remainder of this paper is structured as follows. We describe several types of robotic implementations using various existing heart rate detection approaches in Section 2. Section 3 explains the fundamental methodology and rationale for the hardware selected and details the software design and implementation. A description of our approach to the experiment design and test results can be found in Section 4. We conclude the paper with a summary of our research and a brief discussion of future work in Section 5.

2 RELATED WORK

Remotely detecting shifts in the psychological state and physical condition of humans is a challenge recently undertaken by researchers in numerous fields including image and signal processing [5], human-computer interaction [6], computer vision [7], biomedical engineering [8] and robotics [9]. Until very recently, studies in human-robot interaction have typically obtained physiological information from humans using contact modalities such as wearable biofeedback sensors or sensors fitted on the robot. These initial studies provide valuable insight for understanding how physiological indicators can yield critical information pertaining to the affective state of humans interacting with robots.

2.1 Vital Signs Detection with Interactive Robots

Important research using robots and contact sensors have shown that physiological responses alone can be used to successfully recognize affective states in humans [10–12]. In the first study three physiological indicators — heart rate, skin conductance and facial muscle contraction — are collected to perform affective state estimation during human-robot interactions. While face muscle contraction was not found to be strongly correlated with affective state in the set of tests performed, heart rate acceleration was found to be one of three important physiological features for successfully predicting affective state.

Other research applies the fundamental concept of human stress detection to the study of autism therapy [13]. Participants of the study were fitted with biofeedback sensors which measured heart rate variability, skin conductivity, eyebrow movement, jaw clenching, and body temperature. The key to this approach includes designing an affective control architecture and creating rules by which the robot decides how to respond when the threshold anxiety level is reached.

A study in human-robot interaction collects physiological signals using wearable biofeedback sensors to recognize the affective state of the human and adapt the behavior of the robot accordingly [14]. Experiments were designed using a robot-based basketball game where a robot monitors the participant's anxiety and alters the difficulty level of the game based on the perceived stress level of the player. Results show that overall player performance is improved when the difficulty level is adjusted based on physiological cues and not merely on performance alone.

The modes of collecting physiological data described in these studies have been effective but each method still requires that the subject be fitted with the proper biofeedback sensors. In some controlled settings and with certain subjects, this may not be an issue. However, their efficacy in dynamic environments where people cannot be fitted with biofeedback sensors or in certain medical or therapeutic settings where persons are averse to wearing sensors is still somewhat limited. While the collection of physiological data for diagnosing disorders and stress in humans is not new, remotely recovering this information for use in robotics is an emerging field. This paper presents a complementary approach in which heart rate, an important physiological indicator, is collected using a non-contact modality.

2.2 Remote Detection of Vital Signs

Non-contact modalities have been explored including laser doppler vibrometry (LDV) [15], radio frequency scanners [16] and microwave doppler radar [17]. One study remotely collects physiological information using LDV to deduce the stress state of an individual based on vibrations of the skin directly covering the carotid artery [15]. As currently implemented, this approach yields an intersession equal-error rate of 6.3%. The main drawbacks to these approaches include problems with accurate tracking due to variances in patient physiology and the prohibitive cost of the technology.

The biomedical engineering field has also published a great deal of research dedicated to the acquisition of a wide variety of physiological information. In one study, a low-cost camera is used to detect subtle skin color changes over time in order to deduce heart rate [8]. Video recordings were analyzed using independent component analysis on three color channels (RGB) to extract three important physiological indicators: heart rate, respiratory rate and heart rate variability. The research showed that using a camera alone can yield fairly accurate results. However, there are significant limitations using color as the sole measure of physiological indicators including common variations in skin tone, ambient or direct lighting and proper face detection due to even small changes in illumination, shadows and occlusions.

Studies in the electrical engineering field have also approached the challenge of remote vital signs monitoring by targeting a Doppler radar at a person's chest to measure small changes in the demodulated voltage waveform which represents displacement due to respiration and heart activity [18]. Several signal enhancement techniques are applied including center clipping and a Hanning window before the heart rate is computed. Although this modality provides reasonably accurate results, doppler radar is highly sensitive to motion originating from both the target and from extraneous motion within the general range of the antennae [19]. Our research uses a simpler, more robust approach that relies on temperature, which is tolerant of color and movement artifacts, to detect heart rate and does not require the use of expensive equipment or processing a large amount of complex data.

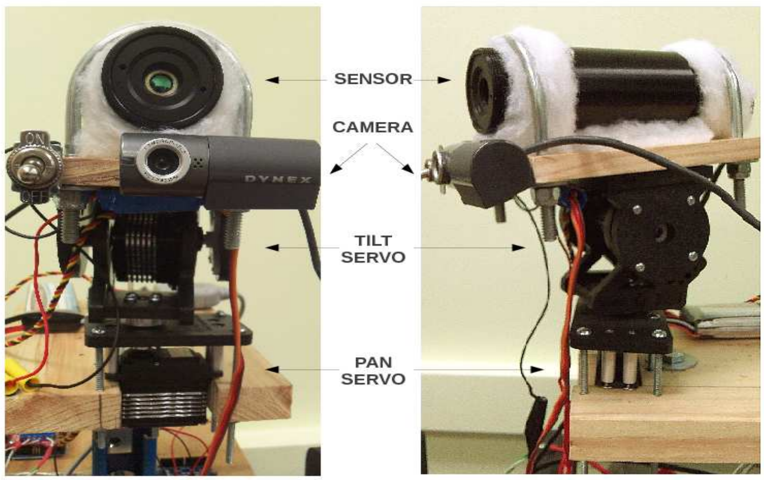

Fig. 1. Remote breathing monitoring system. Front view (left) and profile view (right).

3 METHODOLOGY

This research presents a new technique for remotely detecting and monitoring heart rate in real time. Although resting heart rates may vary from one individual to the next, healthy adults have a typical resting heart rate between 60-120 bpm [20]. Each beat of the heart consists of a series of deflections reflecting the time evolution of electrical activity in the heart that is responsible for initiating muscle contraction. A single heartbeat is typically decomposed into five constituent parts labeled: P, Q, R, S, and T. The largest-amplitude portion of the ECG is the QRS complex, caused by currents generated when the ventricles depolarize prior to their contraction. We are most interested in measuring the QRS component of the cardiac cycle where one heart beat is measured from the beginning of one QRS cycle to the beginning of the next.

The heart rate measurement system described in this research (see Figure 1) employs a single-point infrared sensor introduced in our previous research [4] for collecting and calculating non-contact breathing rates. The process for temperature data collection consists of the following: (1) aim the sensor at a pre-defined sub-nasal target region using the location of the nose as extracted from the most recent video frame and, (2) extract the temperature information provided by the sensor analog signal. The subnasal region is selected as the target for two reasons. First, the superior labial artery, which follows a course along the edge of the upper lip, is believed to cause subtle temperature changes corresponding to cardiac pulse. Second, we seek to extract heart rate and breathing rate concurrently in future applications using the proposed sensor mechanism.

In order to accomplish proper sensor targeting and temperature extraction, a specific combination of hardware and software was included in the overall system design. Our system uses the same custom-built actuated platform enumerated in our previous work [4] so we omit a detailed description here.

3.1 SOFTWARE

The software developed for our system manages three main functions: (1) infrared sensor positioning, (2) temperature collection, (3) data pre-processing and heart rate calculation. Although sensor positioning is accomplished using the same technique described in our prior study [4], new software was developed for an improved rate of temperature collection, a more robust approach for data processing and the extraction of a heart rate. Sensor positioning relies on repeated nose detections and automatic adjustment of the sensor's platform position in order to maintain the subject's nose in the target region of the IR sensor. The IR sensor is repeatedly sampled and collected data is subsequently processed in order to extract a temperature in degrees Fahrenheit. Once the raw temperature data is pre-processed using a low-pass filter, a DWT is used to compute the heart rate in beats per minute. The data processing method implemented for detecting heart rate is sufficiently fast to be used in real time.

Infrared temperature collection The infrared temperature collection system has been improved so that a sample rate of 20-25 samples per second is achieved, compared to the 6 samples per second collected by the original system. This represents the upper limit of the sample rate we can achieve given the hardware employed and is accomplished by processing the nose detection and sensor positioning independently from the temperature collection and data processing. Because heart beat events occur at a much higher frequency than breathing events, increasing the number of samples collected per second was necessary in order to capture the relatively short-lived temperature increases that correlate to heart rate.

The infrared sensor is continuously sampled until a window of 32 time-stamped samples or approximately 1.6 seconds of temperature data has been collected. Various window sizes were tested in order to evaluate the system's performance during periodic heart rate fluctuations. Although larger window sizes provide higher stability in computed heart rates, they are prone to excessive smoothing and reduce the system's ability to detect short-lived heart rate increases or decreases. Further, while small window sizes are susceptible to being dominated by relatively small errors that can be introduced when temperatures are collected during re-targeting, they provide more resilient and responsive heart rate detection overall.

Data pre-processing Data is pre-processed in three steps. First, temperatures that are collected when the infrared sensor is performing initial targeting and periodic re-targeting are often too low to be considered related to the human body. Exceptionally low readings are assumed to be from a non-human source and are excluded from the data set at the time of temperature collection. Secondly, low-amplitude noise picked up by the sensor signal is minimized using a low-pass filter on each collected set of data. Finally, to make the IR data suitable for processing with a DWT, the 0-mean is computed for all the samples in each window.

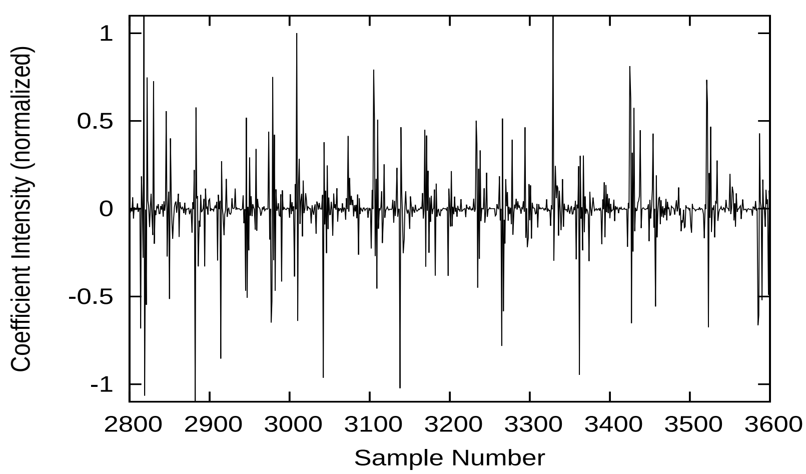

Fig. 2. Sample set of DWT coefficients.

Heart rate calculation Heart rates are computed using a DWT [21] on each window of collected infrared data. A DWT is used to process the IR data for two fundamental reasons: (1) heart rates are not stationary since they have varying frequency components at different time intervals and, (2) we are interested in the temporal information associated with each reading. Unlike Fast Fourier Transforms (FFTs), DWTs are capable of extracting specific frequencies occuring at particular time intervals.

The DWT first sends samples through a low pass filter which yields approximation coefficients and a high pass filter which results in one or more detail coefficients. The outcome of this filtering technique is that the component signal frequencies are cut in half and according to Nyquist's rule, half the samples can be discarded. Although this process halves the time resolution and each output has half of the input frequency band (since only half of each filter output characterizes the entire signal), the frequency resolution is effectively doubled with each decomposition.

The decomposition process is recursively repeated in order to increase the frequency resolution until no further decompositions are possible. Once the decomposition is completed, a set of coefficients is output that were produced at various scales and at different time intervals of the signal. The coefficients can then be analyzed to extract frequency information for particular time intervals or for the signal in its entirety.

Due to the nature of this technique, the number of samples processed in a given data set by the DWT must be in powers of two. Our system uses the Daubechies 6 (db6) [22] wavelet to perform the transform and collects temperature readings at a rate of 20 samples per second so the highest frequency that can be extracted is 10 samples per second or 10 Hertz (Hz). The range of frequencies in which we are most interested for this research are 0.8-1.90 Hz because they correspond to heart rates between 48 bpm and 114 bpm. The DWT levels of decomposition which contain the detail coefficients within that frequency range are found at levels 3 and 4 and represent frequencies between 1.25-2.5 Hz (level 3) and 0.625-1.25 Hz (level 4). Since the input signal is recursively decomposed into component frequencies using a DWT, it is possible to isolate breathing from heart rates, with effectively no crossover. Figure 2 illustrates a representative coefficient file produced for one set of pre-processed IR data.

Finally, the heart rate is extracted by (1) computing and comparing the average amplitude of the detail coefficients at levels 3 and 4, (2) selecting the level with the largest average amplitude, (3) counting the number of zero crossings for the coefficients at the selected level and, (4) multiplying that number by 37.5 (the number of 1.6-second windows in a minute). Zero crossings are defined as any change in signal direction which exceeds a minimal threshold of 0.1.

4 EXPERIMENTS

Experiments were conducted to measure the effectiveness of the single-point infrared sensor for detecting heart rates remotely. A representative graph of extracted heart rates as detected by the IR sensor and by the ECG illustrate typical results over a period of approximately 35 seconds in Figure 3.

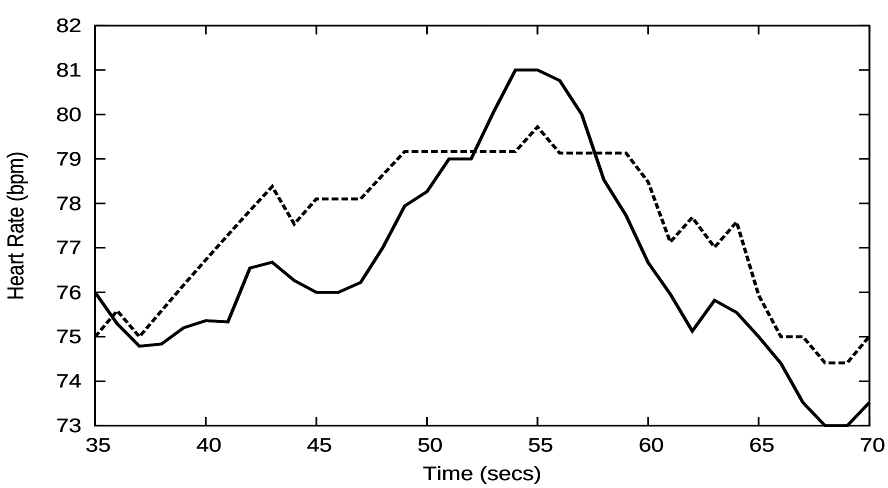

Fig. 3. ECG and IR heart rates for a representative 35-second data set. IR heart rate (dashed line) and ECG heart rate window (solid line).

4.1 Experiment setup

The proposed system's accuracy were quantitatively measured by collecting temperature data with the infrared sensor, computing the heart rate and comparing the results with heart rate data obtained from an ECG. For the ECG data collection, participants were each fitted with 3 electrodes attached to a bioradio which continuously transmitted heart rate data to a nearby computer. ECG information is collected at approximately 600 samples per second and a heart rate is computed for each 960 samples, or 1.6 seconds of ECG data, so that IR and ECG heart rates can be easily processed and compared.

Experiments included 24 study participants, 17 females and 7 males, between the ages of 18 and 35. Individuals who participated in this study were not taking medication which could interfere with their heart rate at the time of the experiment. Each participant was asked to sit in a chair that was situated approximately 1 meter from a rolling table equipped with the infrared sensor system and a laptop computer. During the course of each 10-minute test session study subjects watched a video playing on the laptop computer. The primary purpose of the video was to maintain the participant's attention in a forward-facing, relatively still position. Small movements resulting from participants shifting their position during the test session were automatically managed using incremental pan and tilt adjustments of the sensor platform.

4.2 Experiment Results

Twenty four test sets, each consisting of approximately 10 minutes of data and about 375 individual heart rates were collected and analyzed. Of those 24 sets, two were identified as anomalous due to obvious and persistent nose detection problems observed while the test was being conducted. Common problems in feature detection are typically due to false positive identification of other artifacts in the environment that possess similar characteristics to the target feature. For example, during one test set the nose detection system falsely identified the subject's eye as her nose and the entire test set collected data consisting of temperatures measured around the eye region. Twenty two test sets are classified as "typical" and contain data collected when the nose detection and tracking was not clearly working improperly.

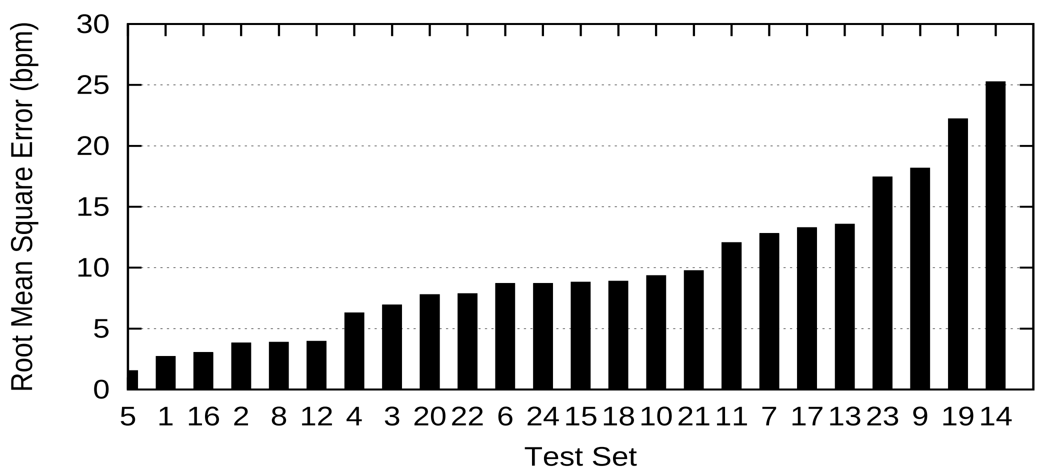

Fig. 4. All 24 test set results by RMSE. Sets 14 and 19 are considered atypical because of obvious nose detection issues during testing.

Overall system accuracy was measured by computing the difference between the reported ECG heart rate and the IR detected heart rate for each 1.6-second window (Figure 4). Because the ECG data collected during experiments consists of a heart rate without a time-stamp, part of the system performance analysis includes an auto-correction for the temporal alignment of data between ECG heart rates and IR-derived heart rates by comparing the root-mean-square errors (RMSE) of various offsets for each window of coefficients computed. Typical and anomalous test sets were analyzed separately and accuracy was assessed in beats per minute (bpm). Six categories were used to classify our results: (1) 0-4 bpm, (2) 5-9 bpm, (3) 10-14 bpm, (4) 15-19 bpm, (5) 20-24 and (6) 25 and higher bpm.

Of the approximately 375 heart rates compared for each typical test set, about 73% (or 16 out of 22) averaged heart rates within 0-9 beats per minute as compared to average heart rates produced by the ECG over the entire 10-minute test set (Table 1).

Range of RMSE in bpm	Percentage of Typical Test Cases
0-4	27.3%
5-9	45.4%
10-14	18.2%
15-19	9.1%

Table 1. Typical test sets results by RMSE. Two excluded sets had persistent nose detection problems.

Test Set	Lowest ECG Heart Rate	Highest ECG Heart Rate	Accuracy in RMSE
1	71	87	2.71
2	61	74	3.81
4	93	115	6.28
16	71	88	3.03

Table 2. Representative test sets successfully detecting lower- and higher-range heart rates.

Although stress testing was not conducted in these initial experiments, heart rates within a given session varied on average by about 20 bpm. The minimum and maximum heart rates detected with low RMSE scores ranged between 61 and 115 bpm. A summary of results demonstrating several successful test sets for both lower range heart rates and higher heart rates is included in Table 2.

An additional consideration in the assessment of system performance is the system's ability to effectively track increases and decreases in heart rate even when the baseline is shifted by an offset as shown in Figure 5. Test sets that mirror heart rate fluctuations as reported by the ECG but are offset by a certain amount will produce higher RMSE scores on average even though increases and decreases in heart rate are accurately detected. We believe there are two potential causes for this offset. First, computing the number of zero crossings exclusively for the captured IR data may not be sufficient to calculate the heart rate accurately. Second, intermittent errors in targeting the infrared sensor precisely may result in cumulative errors in computed heart rate. Future work will include an evaluation of these cases to determine if a baseline shift can be corrected and if they can still be used to provide valuable information pertaining to changes in heart rate that are indicative of stress state.

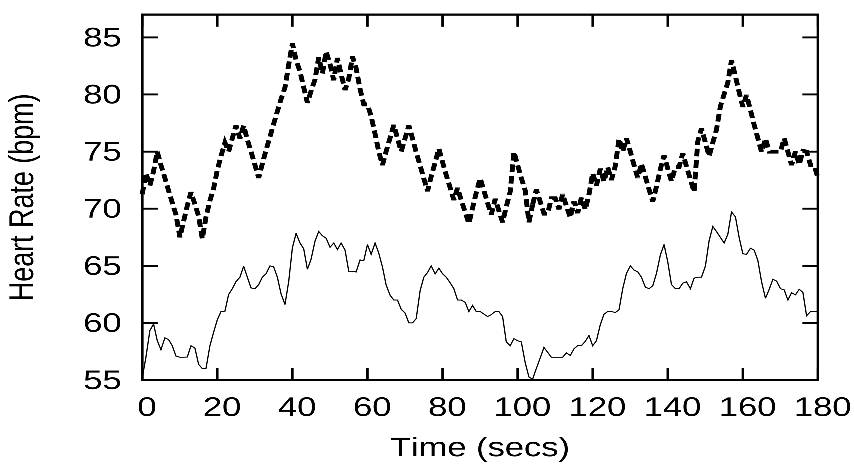

Fig. 5. A 3-minute data excerpt illustrating a baseline offset between IR heart rate (upper line) and ECG heart rate (lower line).

5 CONCLUSIONS

This paper presents a new non-contact technique for measuring temperature changes in the sub-nasal skin surface temperature to calculate heart rate. This study extends and improves our previous work where breathing rates were automatically extracted using a curve-fitting function and ground truth was measured with a self-reporting method. Several enhancements were made to the original system software and testing design, making it possible to use the same hardware to remotely extract a heart rate.

Initial results from the tests conducted in this study are very promising. This study demonstrates a low-cost, potential solution for obtaining physiological information using a non-contact approach. Due to its small size and the minimal computation required for the calculation of heart rates, incorporating such a device in robots working alongside humans in many SAR and RR applications has great potential.

Future work will focus on improving the accuracy and robustness of the sensor targeting component in order to minimize heart rate detection errors resulting from occasional or persistent drifting. Enhancements to the nose detection system will include detecting other facial landmarks to assist the classifier in identifying the nose and specifically, the sub-nassal region, more accurately. Additionally, further testing will be conducted to examine the system's ability to detect heart rates effectively when the user is engaged in light activity and sensor positioning must respond quickly to minimal but frequent user movements.

References

1. H.I. Krebs and N. Hogan. Therapeutic robotics a technology push. *Proceedings of the IEEE*, 94(9):1727–1738, 2006.
2. J.D. Montagu and E.M. Coles. Mechanism and measurement of the galvanic skin response. *Psychological Bulletin*, 65(5):261–279, 1966.
3. P.D. Tyson. Task-related stress and EEG alpha biofeedback. *Biofeedback and Self-Regulation*, 12(2):105–119, 1987.
4. L. Boccanfuso and J.M. O'Kane. Remote measurement of breathing rate in real time using a high precision, single-point infrared temperature sensor. In *Proceedings IEEE International Conference on Biomedical Robotics and Biomechatronics*, 2012.
5. T. Matsukawa and K. Yokoyama. Visualizing physiological information based on 3dcg. *FORMA*, 25:11–14, 2010.
6. A. Barreto, J. Zhai, and M. Adjouadi. Non-intrusive physiological monitoring for automated stress detection in human-computer interaction. In M. Lew, N. Sebe, T. Huang, and E. Bakker, editors, *Human-Computer Interaction*, volume 4796 of *Lecture Notes in Computer Science*, pages 29–38. Springer Berlin / Heidelberg, 2007.
7. I.A. Kakadiaris, G. Passalis, T. Theoharis, G. Toderici, I. Konstantinidis, and N. Murtuza. Multimodal face recognition: Combination of geometry with physiological information. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1022–1029, 2005.
8. M.Z. Poh, D.J. McDuff, and R.W. Picard. Advancements in noncontact, multiparameter physiological measurements using a webcam. *IEEE Transactions on Biomedical Engineering*, 58(1):7–11, 2011.
9. S. Suzuki, T. Matsui, S. Gotoh, Y. Mori, B. Takase, and M. Ishihara. Development of non-contact monitoring system of heart rate variability (hrv) - an approach of remote sensing for ubiquitous technology. In *Proceedings of the International Conference on Ergonomics and Health Aspects of Work with Computers: Held as Part of HCI International*, pages 195–203, 2009.
10. D. Kulic and E. Croft. Estimating robot induced affective state using hidden markov models. In *The IEEE International Symposium on Robot and Human Interactive Communication (ROMAN)*, pages 257–262, Sept 2006.
11. D. Kulic and E. Croft. Physiological and subjective responses to articulated robot motion. *Robotica*, 25(1):13–27, 2007.
12. P. Rani, J. Sims, R. Brackin, and N. Sarkar. Online stress detection using psychophysiological signals for implicit human-robot cooperation. *Robotica*, 20(6):673–685, 2002.
13. L. McLaughlin. Engineering meets emotion: New robots sense human distress. In *Intelligent Systems, IEEE*, volume 18, pages 4–7, 2003.
14. C. Liu, P. Rani, and N. Sarkar. Affective state recognition and adaptation in human-robot interaction: A design approach. *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 3099–3106, 2006.
15. J.W. Rohrbaugh, E.J. Sirevaag, N. Singla, J. O'Sullivan, and M. Chen. Laser doppler vibrometry measures of physiological function: evaluation of biometric capabilities. *IEEE Transactions on Information Forensics and Security*, 5:449–460, 2010.
16. *A mini-robot with a long standoff capability to detect motion and breathing inside a compound*, September 2011. http://www.ereleases.com/pr/author/tialinx-inc (accessed January 2012).
17. M. Sekine and K. Maeno. Non-contact heart rate detection using periodic variation in doppler frequency. In *Proceedings of the IEEE, Sensors Applications Symposium (SAS)*, pages 318 –322, Feb 2011.
18. B. Lohman, O. Boric-Lubecke, V.M. Lubecke, P.W. Ong, and M.M. Sondhi. A digital signal processor for doppler radar sensing of vital signs. *IEEE Engineering in Medicine and Biology Magazine*, 21(5):161–164, 2002.
19. A. Singh and V.M. Lubecke. Respiratory monitoring and clutter rejection using a cw doppler radar with passive rf tags. *IEEE Sensors Journal*, 12(3):558–565, march 2012.
20. P.E. McSharry, G.D. Clifford, L. Tarassenko, and L.A. Smith. A dynamical model for generating synthetic electrocardiogram signals. *IEEE Transactions on Biomedical Engineering*, 50(3):289–294, Mar 2003.
21. CS Burrus, RA Gopinath, and H Guo. *Introduction to Wavelets and Wavelet Transforms: A Primer*. Prentice Hall, 1997.
22. I. Daubechies. *Ten lectures on wavelets*. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 1992.