

Particle Filtering and Sensor Fusion for Robust Heart Rate Monitoring using Wearable Sensors

Viswam Nathan, *IEEE Student Member*, and Roozbeh Jafari, *IEEE Senior Member*

Abstract—This article describes a novel methodology leveraging particle filters for the application of robust heart rate monitoring in the presence of motion artifacts. Motion is a key source of noise that confounds traditional heart rate estimation algorithms for wearable sensors due to the introduction of spurious artifacts in the signals. In contrast to previous particle filtering approaches, we formulate the heart rate itself as the only state to be estimated, and do not rely on multiple specific signal features. Instead, we design observation mechanisms to leverage the known steady, consistent nature of heart rate variations to meet the objective of continuous monitoring of heart rate using wearable sensors. Furthermore, this independence from specific signal features also allows us to fuse information from multiple sensors and signal modalities to further improve estimation accuracy. The signal processing methods described in this work were tested on real motion artifact affected electrocardiogram (ECG) and photoplethysmogram (PPG) data with concurrent accelerometer readings. Results show promising average error rates less than 2 beats per minute (bpm) for data collected during intense running activities. Furthermore, a comparison with contemporary signal processing techniques for the same objective shows how the proposed implementation is also computationally more efficient for comparable performance.

Index Terms—Particle Filter, Physiological Signal Processing, Motion Artifacts, Heart Rate, Wearable Sensors

I. INTRODUCTION

IN recent times, a significant amount of research has been dedicated to realizing the goal of pervasive, around-the-clock personal physiological monitoring that is both easy to use and accurate. The traditional norm is to rely exclusively on hospital visits, however this model of healthcare could be significantly augmented with the help of dedicated wearable and environmental sensors. Apart from generally being more convenient and cheaper, continuous monitoring with a wearable device will be invaluable for identifying medical conditions for which sporadic monitoring is insufficient. Among the various physiological measures, there has been significant interest in wearable continuous heart rate (HR) monitoring, and it would be useful for both fitness and health applications. For example, heart rate variability is a condition

that could be an indicator of myocardial ischemia [1] and continuous monitoring of the heart rate would increase the chances of a successful diagnosis. Wearable heart rate monitors could also enable general fitness applications. For example, users would be able to tailor their exercise routines according to how their body is responding to the workload. Continuous cardiac activity monitoring could also be used for user authentication purposes, as shown in a recent work [2].

However, the increased comfort and convenience of these sensors often comes at the cost of an increased amount of noise as well; for example, for bio-potential sensors, the skin contact would likely have to be a dry electrode, potentially non-contact electrodes, which in turn means an increased amount of noise compared to wet or gel-based interfaces which are still the standard for medical grade equipment [3]. Moreover, activities of daily living are likely to lead to motion artifacts in the signals, such as in the case of a smart watch monitoring heart rate while the user is jogging. A recent study exploring continuous monitoring of heart rate of construction workers to manage workload showed that motion artifacts were the primary source of error in heart rate estimation [4].

There are multiple modalities, sensor types, and sensor locations that can be used to capture heart rate. The heart rate can be inferred from the ECG, PPG or even the changes in bio-impedance of a certain region of the skin. The type of sensors can be gel-based patches, or dry metal electrodes or an LED-photodiode combination as is the case for PPG. Finally, the sensors could be available in multiple locations depending on the context such as on a T-shirt, on a wrist watch, or on the sides of reading glasses. Not to mention the possibility of environmental sensors opportunistically measuring the heart rate, such as a steering wheel capturing ECG activity with electrodes placed on either side of the wheel. Moreover, apart from all these sources for measuring the ‘signal’, we can also easily imagine sources that measure the ‘noise’. An obvious example would be accelerometers that are already part of portable devices such as smart watches; these signals could be correlated with the motion noise that clouds the estimation of heart rate. The internet of things (IoT) would very much include physiological and motion sensors, and it is very likely that successful signal processing techniques would take advantage of the expected multitude of sensors. Therefore, we strive to show that the proposed implementation is capable of supporting multiple simultaneous and heterogeneous sensors and fusing them effectively for more accurate estimates.

Finally, an important consideration when designing

Submitted for review on August 30th, 2017. This work was supported in part by TerraSwarm, one of six centers of STARnet, a Semiconductor Research Corporation program sponsored by MARCO and DARPA.

V. Nathan is with the CSE department and R. Jafari is with the BME, CSE ad ECE departments at Texas A&M University, College Station, Texas, 77843, USA. (e-mails: {viswamnathan,rjafari}@tamu.edu).

wearable sensors is the computational power. Due to form factor and battery lifetime constraints, the processors in these systems are relatively modest, and hence it is important to ensure that the computational load due to signal processing algorithms is minimized as much as possible.

In this work we describe the use of a particle filter to estimate target physiological phenomena from a variety of signal modalities, as well as the fusion of these. The particle filter is a sequential Monte Carlo routine that probabilistically estimates the true state of a given system by updating the weights and redistributing a set of ‘particles’. More details on the working of the particle filter will be provided in this article and can also be found in other works [5].

Moreover, while the focus of this work is on heart rate monitoring, the proposed implementation of the particle filter could also be used for estimating other physiological parameters that are relatively slow-changing and consistent over short time intervals. As will be explained in detail later, based on sensor observations, the particle filter tracks multiple possibilities for the target parameter and rewards those that are consistent over time. We know that the human heart rate is relatively steady over short time intervals, but this applies to other phenomena as well, such as respiration rate or continuous arterial blood pressure (ABP). Continuous ABP can be estimated by measuring the pulse transit time [6], which in turn can be estimated with the same signal modalities discussed in this work. As such, not only is the proposed implementation independent of specific signal modalities or features, it is also potentially adaptable for other applications in the same domain. The contributions of this work are:

1. A particle filter formulation for estimation of heart rate from noisy sensor streams without dependence on specific signal features; it instead works with naïve, greedy observation mechanisms and leverages the expected steady changes of the human heart rate.
2. Demonstration of the efficacy of the technique using real motion-artifact affected ECG and PPG data.
3. Showcasing our implementation’s potential for fusion of multiple signal modalities to improve heart rate estimates.
4. Demonstration of improved computational efficiency of our solution compared to contemporary related works.

II. RELATED WORKS

There have been many proposed approaches in the literature to obtain an accurate heart rate estimate from a noisy cardiac signal. Several works have been based on the use of an adaptive filter, but such techniques always rely on the presence of an external reference signal, such as accelerometer data [7, 8] or electrode tissue impedance [9], which may not always be available. Moreover, different reference signals may be better correlated with different types of motion artifacts and thus a system based on only one reference signal may not represent a generalized solution to handle artifacts from a variety of user actions.

Methods based on a Kalman filter do not rely on an external reference [10], but these techniques assume that the signal and observation models are linear functions and that the noise is

Gaussian, which is not always the case for biomedical applications [11]. The extended Kalman filter was introduced to circumvent the disadvantage of the linearity assumption [12], but just like the regular Kalman filter it still suffers from the fact that only unimodal Gaussian distributions can be tracked [11]. In other words, only one possibility for the true state can be tracked at a time and if the estimate diverges from the true state, it may continue to diverge beyond recovery.

Apart from these, there have been a few works that successfully combine several signal processing techniques along with heuristic knowledge of signal characteristics to build a heart rate detection algorithm. There have been three recent related works of note that tackle the problem of heart rate estimation in the presence of extreme motion artifacts when running. The first, dubbed TROIKA [13], involves primarily singular value decomposition, an optimization approach to find a sparse signal representation of the PPG frequency spectrum and finally spectral peak tracking approaches to estimate the heart rate. The second technique developed by the same author, called JOSS [14], has a similar approach except it jointly estimates the spectra of the accelerometer as well as the PPG and does away with certain steps to save on computation time. The third and final work for PPG signals with motion artifacts [15], which we will refer to as ‘Robust EEMD’, is based on ensemble empirical mode decomposition (EEMD), followed by a recursive least squares (RLS) adaptive filter using the accelerometer signal as reference. These two techniques are followed by several spectral peak tracking approaches as well as heuristic conditional steps to track the heart rate frequency. Later in this paper, we will present a comparison of our proposed particle-filter based approach with these three works, both in terms of estimation accuracy as well as computational efficiency.

The particle filter is a probabilistic method that does not depend on any external reference signal nor assume a specific distribution for either the signal or the noise as is the case for the Kalman filter. It is robust and has the potential to recover from incorrect estimates since it can keep track of multiple possibilities. It is generalizable and can be adapted to handle a variety of signal and noise models. It is also straightforward to adjust the number of particles in use, to trade-off between computation time and accuracy depending on the application.

The particle filter has been previously employed in other similar applications, such as identifying the various segments of an ECG in stationary conditions. However, apart from not dealing with motion artifacts, these works usually incorporate a complex dynamical model for the ECG that involves several state dimensions, which in turn increases the computational cost [16, 17]. Another work based on an ECG model has a much reduced dimensionality for the state space; however it is only tested for ECG contaminated by white or pink noise [18].

A particle filter has also been employed for muscle artifact affected ECG de-noising, however this also relies on a sophisticated model that is specific to the progression of ECG with multi-dimensional states and does not seem to be validated on ECG signals with a significant amount of noise [19]. Moreover, in all of the above, the approach that relies on

the use of a single rigid and specific mathematical model may not be generalizable to be used for a wider variety of signals from different subjects [20]. The key difference in our proposed framework is that the heart rate itself is directly used as the state to be estimated in the particle filter model equations, and we design the observation densities such that the particle filter simply rewards those observations that are consistent with the expected behavior of a human heart rate. Moreover, the use of only a single state dimension in the formulation greatly eases the computational load compared to previous particle filter implementations in this domain.

III. BACKGROUND

A. Particle Filter

In order to formulate the state estimation problem for heart rate detection, we first define the state space representation:

$$\mathcal{X}_t \sim \pi_x(\mathcal{X}_t) \quad (\text{initial distribution})$$

$$Z_t | \mathcal{X}_t \sim g(\mathcal{X}_t) \quad (\text{observation density})$$

$$\mathcal{X}_{t+1} | \mathcal{X}_t \sim f(\mathcal{X}_t) \quad (\text{transition density})$$

where \mathcal{X}_t denotes the true system state, *i.e.*, the true heart rate at time t , $\pi_x(\mathcal{X}_t)$ denotes the initial distribution of the system states based on some prior knowledge, Z_t denotes a set of discrete *observations*, $g(\cdot)$ is a function representing the observations conditioned on the true heart rate, and $f(\cdot)$ is the state dynamics or transition model that characterizes the heart rate dynamics as a function of time. In essence, the function $g(\mathcal{X}_t)$ denotes the likelihood of observations given the true state, and the function $f(\mathcal{X}_t)$ describes the progression of the true state due to its own dynamics over time.

The state estimation problem can be delegated to a particle filter, which is a sequential Monte Carlo method that solves the problem by maintaining a set of weighted particles, each being a candidate state estimate, its weight being proportional to the likelihood of that particle representing the true state. At each step of the particle filtering problem, the goal is to estimate the posterior state distribution ($p(\mathcal{X}_t | Z_t)$), *i.e.*, the probability distribution of the current true state given a set of observations. This is estimated by the weighted sum:

$$p(\mathcal{X}_t | Z_t) = \sum_{p=1}^{N_p} W_{x_t^p} \delta(\mathcal{X}_t - \mathcal{X}_t^p) \quad (1)$$

\mathcal{X}_t^p is the p^{th} particle at window t ,

$W_{x_t^p}$ denotes the weight of particle \mathcal{X}_t^p ,

N_p is the total number of particles and

$\delta(\cdot)$ is the Dirac delta function, used to place a mass at the particle's location in the posterior probability density function.

Once this posterior probability distribution is updated in each time instance, a suitable method can be used to best estimate the target state at each time. In this work, we use the *maximum a posteriori* (MAP) estimate.

B. Problem Characteristics

There are a number of reasons why the particle filter is a good fit for the particular problem of heart rate estimation in noisy signals, when compared to other similar techniques. For

example, if we consider the problem of heart rate estimation using peak detection on the ECG, a common source of noise is motion that causes spike-like artifacts. These could lead to false positives for a peak detection algorithm. If we consider a specific instance with a true heart rate of 60bpm, if there is a false positive peak between two true peaks, then the average estimated heart rate for that period becomes, say 120bpm. Thus, it is clear that the noise cannot be modeled as a Gaussian distribution around the true value. The probability distribution of the heart rate is in fact multi-modal with several distinct possible heart rates in the probability space. This is precisely why, as mentioned earlier, it may be unsuitable to use the Kalman filter which assumes linear Gaussian models, and the Extended Kalman filter that can track only one of these multiple possible modes. Moreover, given this multi-modal probability distribution space where the different modes can be very far apart in terms of heart rate, we decided to minimize the average error by taking the MAP estimate.

The key insight is that the human heart rate is typically a steady, consistent signal over short time windows; in the following sections we will describe the observation mechanisms that allow the particle filter to essentially become a structure that amasses particles in state space regions that show more consistency. In the dimension of time, since the current distribution of the particle filter depends on previous distributions, there is an inherent sense of ‘memory’ to facilitate rewarding of consistency. In the dimension of state space, N different particles can track N different possible heart rates, thus allowing a parallel search for consistency which reduces the chances of permanently going off track.

Another salient point to note is that the particle filter, as implemented in this work, is decoupled from the signal characteristics. In other words, the particle filter simply receives noisy observations of heart rate, but is agnostic to how these observations were obtained and to what signal modalities and features were used. Thus, the particle filter is not married to the particular observation mechanisms described in this work, and any changes to these mechanisms – for example to add sophistication, or make it more suitable for the given sensor or application scenario – can be easily integrated into the same particle filter framework. More importantly, other signal modalities for heart rate detection, such as the ballistocardiogram (BCG), seismocardiogram (SCG) or bio-impedance, could also fit into the same framework and be fused with estimates from existing sensors.

C. Signal Characteristics

Electrocardiogram Signal:

The ECG is a representation of the electrical activity of the heart. In this work we are interested only in the heart rate, and one of the most common ways to estimate the heart rate from ECG is using the R-peaks. The R-peak denotes the point of electrical depolarization of the ventricles of the heart at the start of each beat. The time between successive R-peaks can be used to calculate the beat-to-beat heart rate (Figure 1). However, as mentioned before, ‘spike’-like effects caused by motion artifacts could be falsely identified as R-peaks. This

naturally leads to overestimation of the heart rate when using peak detection based approaches.




Figure 1 - R-peak and R-R interval in an ECG waveform

Photoplethysmogram Signal:

The PPG is obtained by transmitting light of suitable wavelength into the skin and using a photodiode to capture the reflected response that is modulated by the flow of blood, as shown in Figure 2. Unlike ECG, since there is no clear time domain feature, we instead use frequency domain observations for the PPG as will be described in Section IV A.




Figure 2 – PPG waveforms showing periodic heartbeat

Accelerometer Signal

It is quite probable that frequencies due to the cadence of walking or running motions would be prominently present in both ECG and PPG signals. Since these frequencies would also be present in the accelerometer's data, we can leverage this to better inform the estimation process. This is particularly important in instances where the motion results in a 'periodic noise' in either the ECG or PPG signal. The particle filter is designed to exploit the assumed quasi-periodicity of the heartbeat and randomness of motion artifacts; thus in the specific instance of noise due to periodic motion, the accelerometer observations can prove critical to distinguish this from periodic heart rate.

IV. METHODS

A. Observation Mechanism

The sensors provide observations of the heart rate, and the particles update their weights according to each of these.

Photoplethysmogram Observations

The PPG signal is first bandpass filtered between 0.5 and 15Hz to remove baseline wander and unrelated high frequency noise. Subsequently, we use a spectrogram based approach, taking moving, overlapping windows of the PPG stream and applying the short-time Fourier transform. The window size was set to be 8 seconds, with an overlap of 2 seconds between successive windows. The frequency spectrum from a window of PPG constitutes one set of observations.

Electrocardiogram Observations

When processing the time domain ECG signal, we use two back-to-back non-overlapping windows dubbed W_{start} and W_{end} . For the purposes of calculating heart rate, we only consider peak-to-peak intervals that begin with a peak in W_{start} and end with a peak in W_{end} . All such intervals taken together constitute a set of observations for a given time window. Note that the peaks that constitute these pairs may or may not be

artifacts caused by motion or other noise sources. In the example in Figure 3, four different heart rate observations will be considered based on the peak-to-peak pairs shown.




Figure 3 - Windowing illustration on ECG signal

We use windows of size 2 seconds in this work, with a step size of approximately 0.27 seconds. This step size was chosen to accommodate the fact that we expect heart rates as high as 220 beats per minute, and a step size bigger than this could potentially mean skipping true peak-peak observations in those scenarios. The peak detection is then done as follows:

$$findpeaks(W_i, A_{min}, T_{min}) \rightarrow \{P_1, P_2, \dots, P_k\} = P_{W_i} \quad (2)$$

$findpeaks(W_i, A_{min}, T_{min})$ finds the time of occurrence of all peaks in the signal W_i that have amplitude at least A_{min} and such that no two peaks are within T_{min} time of each other, P_{W_i} is the set of peak locations in time, $\{P_1, P_2, \dots, P_k\}$, returned by the 'findpeaks' function. This function is the default implementation found in MATLAB.

This peak detection on its own however is somewhat naïve, so we add an additional step in the procedure for ECG to reduce the number of false positives. We used the continuous wavelet transform (CWT) on the ECG signal with the Mexican Hat wavelet, a center frequency of 0.25Hz and a scale of 5.29 as suggested by a previous work [21]. This helped to accentuate peaks that more closely resemble an R-peak and diminish other trivial peaks. The step described in Equation (2) is then performed on the wavelet transformed signal to obtain the peak locations. It must be noted that this merely reduces the number of false positives but does not eliminate them. Peak-detection based heart rate estimation based solely on the CWT estimate would still overestimate due to false positives, as shown in our previous work [22].

The heart rate observations are then obtained as follows:

$$P_{W_{start}} = \{P_1, P_2, \dots, P_k\} \quad (3)$$

$$P_{W_{end}} = \{P_1, P_2, \dots, P_m\} \quad (4)$$

$$PP_t \triangleq \text{Set of all } (P_b - P_a)$$

$$\forall P_a \in P_{W_{start}}, P_b \in P_{W_{end}}$$

$$Z_t^n = \left(\frac{fs}{PP_t(n)} \right) \times 60 \quad (5)$$

$P_{W_{start}}$ and $P_{W_{end}}$ refer to the sets of peak locations in a starting and ending window respectively, fs is the sampling rate,

PP_t is the set of peak-to-peak intervals with the first peak in a starting window and second peak in an ending window, Z_t^n is the n^{th} heart rate observation of window t expressed in beats per minute (bpm).

Taking all Z_t^n in a given window corresponds to the set of observations Z_t referred to in Section III A when describing

the particle filter's observation model. It must be noted that we take steps to avoid duplicate observations, *i.e.*, preventing the same two peaks taken as a pair in multiple time windows. We also take steps to ensure any observation included in the set is consistent with other observations of similar heart rate from the same time window already in the set; for example, when we have multiple false peaks we could very well have 4 observations of 50bpm within a 2 second window, but it is clearly impossible in reality for all these observations to be true. So in this example, only those pairs of peaks corresponding to 50bpm that are consistent with each other are taken into the set. This step is necessary to avoid an undue preponderance of lower heart rate observations just because of the nature of our relatively naïve observation mechanism.

Accelerometer Observations

The accelerometer data (denoted as ACC) is processed using the same spectrogram approach used for the PPG signal, with identical windowing procedures. Since there are 3 axes on the accelerometer, we strived to combine them into a single spectrogram to provide a unified source of observation for the motion noise over time. Using only one axis on its own was not an option because there was no certainty about which axis captured the most activity across the different subjects in the database. This is presumably due to variations in sensor placement and running styles among the different subjects.

We computed the spectrogram for each of the 3 axes, and then stitched together a combined spectrogram that always included only the maximum of the three available powers for each of the frequencies in each time window. This greedy approach allows to always capture the motion frequencies without unduly diminishing their relative power.

B. Particle Filter Implementation

The initial distribution for the particle filter, $\pi_x(\mathcal{X}_t)$, is defined as follows:

$$\mathcal{X}_t \sim \pi_x(\mathcal{X}_t) = U(HR_{min}, HR_{max}) \quad (6)$$

Where $U(\cdot)$ denotes a uniform distribution between HR_{min} and HR_{max} , the assumed lower and upper limits of the heart rate defined by reasonable human physiological bounds. We made the initial distribution uniform since we have no prior knowledge on the initial heart rate, other than extreme limits.

For the PPG, the probability of an observation with respect to a given state of heart rate is computed as follows:

$$\varphi_t^i = S_t^i / \sum_{n=1}^F S_t^n, \forall i \in (1, F) \quad (7)$$

$$p(Z_t | \mathcal{X}_t) = g(\mathcal{X}_t) = \varphi_t^d \quad (8)$$

S_t^i is the i^{th} element of the vector of observed power spectrum amplitudes (measured as described in Section IV A) in time window t for the PPG signal,

F is the total number of frequencies under consideration,

φ_t^i is the i^{th} element of φ_t , the probability density function that results from normalizing the values of the observed power spectrum to be between 0 and 1 in time window t ,

d , refers to the frequency in the power spectrum that is closest to the heart rate \mathcal{X}_t .

φ_t^d is the probability of the event that the corresponding frequency represents the true heart rate.

This formulation is based on the assumption that a higher power at a given frequency means the more likely it is that that frequency represents the heart rate. However, we know

that with motion artifacts there could be high power at certain frequencies as a result of the cadence of motion. This is where the observations from the accelerometer sensor come in; we formulate the accelerometer observation function such that we reduce the likelihood of a given frequency representing the true heart rate if it is present in high power in the accelerometer power spectrum. The formulation is as follows:

$$\tilde{\varphi}_t^i = \sum_{i=1}^{i+1} \tilde{S}_t^i / \sum_{n=1}^F \tilde{S}_t^n, \forall i \in (1, F) \quad (9)$$

$$p(Z_t | \mathcal{X}_t) = g(\mathcal{X}_t) = (1 - \tilde{\varphi}_t^d) \quad (10)$$

\tilde{S}_t^i is the i^{th} element of the vector of the observed power spectrum amplitudes in time window t of the accelerometer spectrogram,

F is the total number of frequencies under consideration,

$\tilde{\varphi}_t^i$ is the i^{th} element of $\tilde{\varphi}_t$, the probability density function for motion noise that results from normalizing the values of the observed accelerometer power spectrum to be between 0 and 1 in time window t ,

d , is the index of the power spectrum corresponding to the frequency that most closely matches the heart rate \mathcal{X}_t .

$\tilde{\varphi}_t^d$ is the probability of the event that the corresponding frequency is not the heart rate, which for our purposes means it is noise.

For the ECG, in order to create a continuous probability distribution out of the discrete observations, we fit Gaussian distributions around each of the observations resulting in a Gaussian mixture. The probability of a set of observations is then computed as follows:

$$\begin{aligned} p(Z_t | \mathcal{X}_t) &= g(\mathcal{X}_t) = \sum_{n=1}^{o_t} p(Z_t^n | \mathcal{X}_t) \\ &= \sum_{n=1}^{o_t} N(Z_t^n, \mathcal{X}_t, \sigma_z) \end{aligned} \quad (11)$$

Z_t^n refers to the n^{th} heart rate observation in window t ,

O_t is the total number of observations in window t .

$N(Z_t^n, \mathcal{X}_t, \sigma_z)$ denotes a Gaussian distribution with mean equal to the heart rate \mathcal{X}_t in window t , and standard deviation σ_z reflecting the maximum tolerable deviation between the true heart rate and the observation, evaluated at Z_t^n . σ_z is heuristically set to be 3bpm in this work, to ensure that a given particle is reasonably close to an observation to gain weight. Making this parameter too high would mean even unrelated particles gain weight from a given observation, whereas making it too low would too strictly require particles to exactly match the observation to gain weight.

The particle filter is initialized as follows:

$$X_0^p = U(HR_{min}, HR_{max}) \quad (12)$$

$$W_{X_0^p} = \frac{1}{N_p} \quad (13)$$

$$\forall p \in (1, N_p)$$

X_0^p is the p^{th} particle sampled from the uniform distribution between HR_{min} and HR_{max} , defined to be 40 and 220 bpm respectively for this work, at time $t = 0$,

$W_{X_0^p}$ is the initial weight of particle p at time $t = 0$.

N_p is the total number of particles, set to be 300 in this work. Choosing the number of particles affects a trade-off between estimation accuracy and computation time, which we will elaborate further on in Section VI F.

After this initialization, with each succeeding time window, the particle weights are updated as shown in (14). Note that we use the so-called ‘bootstrap filter’ wherein the state transition density is used as the importance distribution, making the weights of the particles directly proportional to the observation density [5]. We chose to do this to simplify the computational load considering the application domain.

$$W_{x_t^p} = p(Z_t | X_t^p) = \begin{cases} \sum_{n=1}^{O_t} N(Z_t^n, X_t^p, \sigma_z), & \text{for ECG} \\ \varphi_t^d, & \text{for PPG} \\ (1 - \tilde{\varphi}_t^d), & \text{for ACC} \end{cases} \quad \forall p \in (1, N_p) \quad (14)$$

X_t^p is the p^{th} particle of window t ,

$W_{x_t^p}$ is the weight of particle X_t^p ,

$N(Z_t^n, X_t^p, \sigma_z)$ is the value of a Gaussian distribution with mean X_t^p and standard deviation σ_z evaluated at Z_t^n ,

φ_t^d is the probability of the event that the frequency corresponding to X_t^p represents the true heart rate.

$\tilde{\varphi}_t^d$ is the probability of the event that the frequency corresponding to X_t^p is not the heart rate, which for our purposes means it is noise.

The weights are all then normalized to be between 0 and 1:

$$\hat{W}_{x_t^p} = W_{x_t^p} / \sum_{r=1}^{N_p} W_{x_t^r} \quad \forall p \in (1, N_p) \quad (15)$$

$\hat{W}_{x_t^p}$ is the weight of the p^{th} particle of window t normalized so the weights form a probability mass function.

Once the particle weights are calculated the well-known sampling importance resampling (SIR) procedure is employed to prevent particle degeneracy [5]:

$$M_t^p = \sum_{r=1}^{N_p} \hat{W}_{x_t^r}, \forall p \in (1, N_p) \quad (16)$$

$$u = \underset{a}{\text{argmin}} |R_U \sim U(0,1) \leq M_t^a| \quad (17)$$

$$X_t'^p = X_t^u, \forall p \in (1, N_p) \quad (18)$$

M_t^p is the p^{th} element of a cumulative sum vector of the normalized particle weights

$X_t'^p$ is the updated state of the p^{th} particle of window t after resampling, and

R_U is a randomly selected number from the uniform distribution between 0 and 1.

After this step, the distribution of particles approximates the posterior probability distribution of the true heart rate state. To get an estimate for the heart rate in the current time window, as mentioned before, we use the MAP estimate. Since the particle weights are now equalized, we instead look to the distribution of particles to capture the most likely estimate.

We cluster the particles belonging to a similar heart rate together, and can say that the largest cluster represents the most likely state as it is analogous to taking the highest weight particle without the SIR procedure. The clusters and the heart rate estimate are thus calculated as follows:

$$C_n \triangleq \text{Set of all } X_t'^m \mid |X_t'^m - X_t'^n| < CS \\ \forall m \in (1, N_p), \forall n \in (1, N_p) \\ E_t = \sum_i C_{max}^i / |C_{max}| \quad (19)$$

C_n is the n^{th} cluster of particles

CS is the maximum spread of a cluster (set to be 3 bpm)

C_{max}^i refers to the i^{th} member of the largest cluster C_{max} , and E_t is the estimate for time window t . For this specific application, the estimate is heart rate in bpm.

The final step in a given iteration of the particle filter is the model-based update that reflects the state transition model defined earlier in Section III. Essentially, as time progress the true heart rate is expected to be dynamic to an extent, and not remain constant. Therefore, the particles are updated accordingly at the end of each time window to approximate this behavior. We assume that the model governing the human heart rate changes over time is a normal distribution:

$$X_{t+1}^p \sim f(X_t^p) \sim N(X_t'^p, \sigma_x) = X_t'^p + (\sigma_x \times R_N \sim N(0,1)) \quad (20)$$

$$\forall p \in (1, N_p)$$

R_N is a randomly generated number from the standard normal distribution, and

σ_x is the standard deviation capturing the expected change in heart rate from one window to the next. With the window step size being 2 seconds, σ_x is heuristically set to be 6 bpm.

The window then shifts to a new section of the signal and the particle filter continues to track the heart rate in this manner iteratively over successive windows.

C. Particle Weighting Assumptions

It can be seen from the formulation for ECG that for a given set of observations in one time window, we consider all the observations as equally likely. We deemed it more generalizable to not rely on any specific features among a set of observations to differentiate them. Instead, we assume that the true heart rate for the subject would make relatively smooth, continuous and gradual changes over time. Leading on from this, we also assume that the observed heart rates as a result of false positive peaks are more random and inconsistent. With these assumptions, our expectation is that even though all observed heart rates are considered equally likely, the particles will build over the correct heart rate as that is observed more consistently over successive time windows.

D. Fusion Technique

Since we have formulated the particle filter with only the heart rate as the state to be estimated, multiple signal modalities and their observation mechanisms can be fused in the same framework. The particle weighting for an arbitrary number of observation sources, *i.e.*, sensors, is given by:

$$W_{x_t^p}^{fusion} = \prod_{s=1}^S p(Z_t^s | X_t^p) \quad (21)$$

$W_{x_t^p}^{fusion}$ is the weight assigned to particle X_t^p when fusing the information from multiple sources of observation

S is the total number of observation sources

Z_t^s is the set of observations in time window t from source s

In essence we assume that since the different sources are observing the same target phenomenon, particles corresponding to states that are observed with higher weight across multiple sources should be rewarded. Conversely, it is unlikely that the same false state would be observed with high probability across multiple sources. In other words, it would be rare for a source of noise to affect sensors with different modalities placed in different locations in the same way.

In this work, for the fusion of ECG, PPG and ACC sensors, the particle weighting process is modified as follows:

$$W_{x_t^p}^{fusion} = p(Z_t^{ECG}|X_t^p) \times p(Z_t^{PPG}|X_t^p) \times p(Z_t^{ACC}|X_t^p) \quad (22)$$

$$p(Z_t^{ECG}|X_t^p) = \sum_{n=1}^{O_t} N(Z_t^n, X_t^p, \sigma_z) \quad (23)$$

$$p(Z_t^{PPG}|X_t^p) = \phi_t^d \quad (24)$$

$$p(Z_t^{ACC}|X_t^p) = (1 - \tilde{\phi}_t^d) \quad (25)$$

Similarly, in the database to be described in more detail in Section V, there are two separate PPG sensors in addition to the accelerometer in a watch-like device; so the formulation above is modified by simply replacing the observations from ECG with the observations from the second PPG sensor.

With this formulation, we can also get an idea of the contribution of each sensor or signal modality to the overall particle filter estimate in each time window, as shown below:

$$\beta_t^s = \sum_{x_t^p \in C_{max}} p(Z_t^s|x_t^p), \forall s \in S \quad (26)$$

β_t^s is the contribution of sensor s to the particle filter estimate in time window t

X_t^p is a particle in the maximum clique C_{max} for window t

Z_t^s is the set of observations in time window t from sensor s

S is the set of all sensors or signal modalities

This contribution can then be normalized with respect to all the sensors in the system and expressed as a percentage:

$$\beta'_t = (\beta_t^s / \sum_{m \in S} \beta_t^m) \times 100 \quad (27)$$

Sensors producing random, noisy observations will likely have a low contribution to the overall particle filter estimate, thus potentially informing dynamic adjustments to the contributions of individual sensors based on perceived signal quality in real time. Moreover, prior knowledge of the increased reliability of one sensor could allow increased weightage of observations originating from that sensor. In this initial work however, we keep it simple and do not assume that any one signal sensor is inherently better than the other. The advantage of this overall method of fusion is that it is simple and generalizable and can easily be reused for different applications as well as an arbitrary number of sensors.

E. Additional Improvements

While the particle filter framework is complete with the above implementation, we found during the course of our experiments with the data that we could make additional improvements to the algorithm to further reduce error for this specific scenario of estimating HR for a running subject:

Hard Thresholding of ACC

We assume that the power of a frequency in the accelerometer spectrum is directly proportional to the probability of that frequency representing motion noise. However, in a few subjects' data there was a harmonic of the movement frequency that was somewhat low in strength but still high enough to mislead the particle filter. Therefore, we modified the ACC probability function to remove from consideration an observation if the power of the corresponding frequency was greater than 10% of the maximum power observed in the accelerometer for that time window.

Detecting ACC Overlap with Heart Rate Frequency

There were a few instances wherein the dominant ACC frequency happened to overlap with the true heart rate frequency. This would be especially problematic with the hard thresholding introduced above. Therefore, we implemented a rough frequency margin around the expected heart rate, and if the ACC frequency under consideration was within this zone, we did not perform the thresholding. This ensured that we did not effectively remove from consideration particles corresponding to the heart rate simply because the ACC frequency was close. The bounds for the margin were set by taking the average of the previous 3 heart rate estimates in Hz and making a conservative bound of ± 0.1 Hz. This corresponds to an assumption that the heart rate would not change by more than 6 bpm in either direction from one time window to the next.

Detecting Resting State

In all of the data, the subject starts at rest at least for a few seconds before beginning any activities. It makes little sense to include the accelerometer observations in these states. Therefore, we first find the magnitude of acceleration in each time window as follows:

$$\tau = \sqrt{(a_x(t))^2 + (a_y(t))^2 + (a_z(t))^2} \quad (28)$$

$a_x(t)$ is the x-axis acceleration for the given time window

$a_y(t)$ is the y-axis acceleration for the given time window

$a_z(t)$ is the z-axis acceleration for the given time window

The observations of the accelerometer are taken into consideration for the final heart rate estimate only if this magnitude was above a certain threshold. The threshold was heuristically determined to be 1.04 g by examining the data. This parameter has to be heuristically set this way in the absence of more sophisticated activity detection algorithms.

Changing Model for Ramping Up of Heart Rate

When the subject transitions from a resting state to walking or running, there is naturally a sudden increase in heart rate in response to the increased workload. There were a few instances where the particle filter was slow to catch up simply because there happened to be observations corresponding to the slower resting state heart rate which turned out to be false observations. In these situations, there is an error in heart rate for a few time windows because the particle filter already had a preponderance of particles around the resting heart rate and continued to see observations consistent with that heart rate. So in a sense there may be some 'latency' for the particle filter estimates to catch up to the true heart rate when there is an abrupt change in the dynamics.

Therefore we wanted to introduce the notion of context-awareness and have multiple operating modes for the particle filter. When we have the accelerometer, we have an independent source of information that provides additional context for the user's current state. When the subject is at rest or running steadily, we do not expect rapid changes in heart rate and so the particle filter model state update (described in equation (20)) will be conservative. Conversely, when the subject's activity level increases rapidly, we can accordingly adjust the model for state update to temporarily allow for

greater changes. A similar idea for this adaptive changing of model equations based on the current context has been previously implemented in other application areas [23].

For our problem, when the subject was previously at rest (as determined by the magnitude threshold) and the ACC magnitude from (28) changes by a significant margin from one time window to the next, we can assume that increased activity has begun. Then for the next few time windows, instead of using the state update equation described in (20), we use the following:

$$X_{t+1}^p \sim f(X_t^p) \sim N(X_t^p, \sigma_x) = X_t^p + (R_N \sim N(\alpha, \sigma_x)) \quad (29)$$

$$\forall p \in (1, N_p)$$

R_N is a randomly generated number from the normal distribution with mean α and standard deviation σ_x . α is a positive bias meant to indicate that on average, the heart rate is expected to increase. It is set to 6bpm in this work. σ_x is the standard deviation capturing the change in heart rate from one window to the next. It is set to be 10bpm, a larger number to reflect the possible rapid changes in heart rate.

The threshold for required change in acceleration magnitude is set to be 0.04 g, and when such a change occurs the alternate update equation (29) is used for a period of 5 time windows. Note that we do not assume the heart rate definitely must increase whenever higher ACC activity is detected, as that is placing too much trust in a rudimentary activity detection approach. Rather, we simply allow the particles to be more spread out than usual for a few time windows when we detect a possible sign of volatility in the heart rate. In other words, there will be more particles than usual in the higher heart rate regions in anticipation of a sudden increase, but there will continue to be particles corresponding to the previous steady state, lower heart rates, and everything in between. This is one of the core advantages of the particle filter, wherein particles can track multiple possible states in parallel. With this context-aware mode-switching approach, the instances of particle filter estimate latency due to sharp heart rate changes was reduced. We did not observe this latency effect for longer than 3 time windows or 6 seconds across all subjects tested in this work, and there was no latency at all for many subjects.

V. EXPERIMENTAL SETUP

A. PPG Database

Motion artifact affected PPG data was taken from the database used as part of the 2015 IEEE Signal Processing Cup (SP Cup) [13]. This data was recorded at a sampling rate of 125Hz using a wrist-worn dual PPG sensor (*i.e.*, two simultaneous channels of PPG) from 12 subjects. The sensor also included a 3-axis accelerometer. Each trial for a subject consisted of 30 seconds of resting, followed by four stages of activity each for 1 minute, and finally 30 seconds of rest again. The four periods of activity consisted of alternating between relatively slower (6km/h or 8km/h) and faster (12km/h or 15km/h) treadmill speeds. ECG was also simultaneously recorded from the chest using wet sensors, and this is used to obtain the ground truth heart rate. Since the three related works mentioned earlier – TROIKA, JOSS and Robust EEMD

– also all worked on the same dataset, we can directly compare the average errors in heart rate estimation.

At this juncture, we note that the JOSS work resampled the data to 25Hz (presumably to ease the computational load) and also truncated the data for 6 of the 12 subjects; this is because that algorithm is entirely dependent on a clean start for the tracking, and half the subjects had signals with some noise to varying degrees even at the initial resting stage. Therefore, to compare with JOSS we also perform the resampling as well as the truncations described in that paper [14]. However, one of the key advantages of the particle filter is the increased ability to recover from going off-track due to the presence of multiple different particles in the state space. So we will also present the results from the un-truncated datasets and show that the particle filter effectively recovers from these ‘false starts’.

B. ECG Database with Simulated Noise

As the dataset described in the previous section had only clean ECG, we wanted to find another solution to obtain motion artifact affected ECG to test the particle filter on the estimation of heart rate from the fusion of simultaneous noisy ECG and noisy PPG. To the best of our knowledge there was no existing database that provided simultaneously recorded ECG and PPG data that were affected by real motion artifacts and also had the ground truth heart rate available.

Therefore, we turned to the MIT-BIH Noise Stress Test Database to get real motion artifact noise and add it to the existing clean ECG signals from the aforementioned Signal Processing Cup database [24, 25]. The MIT-BIH database, including the techniques to synthetically introduce realistic motion artifact noise, is well respected and has been used in several previous works. The owners of the database themselves provided a technique to add calibrated amounts of motion noise data to any given ECG record from their own database such that the desired SNR level is obtained. We have simply adapted this approach to inject the motion artifact noise into the ECG data from the SP Cup database.

In order to test the fusion approach, we used the ECG data injected with motion artifacts in conjunction with the PPG data that is already present in the same database with real artifacts due to the running activity. The particle filter estimates from these fused observations are compared to the heart rate from the unaffected clean ECG data in the database. We ensure that the added motion artifact noise for ECG is proportionally increased in intensity as and when the running speed increases in a given data record. We chose SNR levels of 3dB and -3dB respectively for the slower and faster speeds.

Generating this noisy ECG allows us to illustrate how the particle filter can fuse multiple modalities to improve heart rate estimates compared to using individual sensors.

C. Experimental Data Collection

Even though we believe the methodology of adding noise to the ECG described in the previous section is sound, we readily concede that the ideal scenario would be to have simultaneously collected ECG and PPG data that were both affected by real motion artifacts during the course of the data

collection. Since such a database is lacking in the literature to our knowledge, we conducted a limited data collection of our own to bolster our experimental conclusions. We used a previously developed system called the BioWatch [26] that collected one channel of PPG signals from the wrist and also included an accelerometer. For the ECG, we used a custom platform based on the TI ADS1299, an analog front end for bio-potential signals. One channel of ECG was recorded from the chest using adhesive gel electrodes in a Lead II configuration to be used as the ground truth. A second channel was recorded using a dry electrode that was secured to the forearm just above the BioWatch using medical tape. This was meant to provide an ECG signal that was more susceptible to motion artifacts. Both devices sampled data at a rate of 125Hz and transmitted data to a PC wirelessly using Bluetooth. Data was collected from 5 subjects running on a treadmill after informed consent and protocol approval by the IRB at Texas A&M University (IRB2016-0193D). The experimental protocol was designed to be similar to that of the database described earlier: 30 seconds of rest, followed by four 1-minute periods alternating between walking and running, and 30 seconds of rest at the end. Examples of the signals from our system after pre-processing (0.5 to 15Hz bandpass for PPG and 0.5 to 30Hz bandpass for ECG), for both standing and running scenarios are shown in Figures 4 and 5 respectively.




Figure 4 – ECG, PPG and Accelerometer signals with subject at rest




Figure 5 – ECG, PPG and Accelerometer signals with subject running

VI. RESULTS

A. Heart Rate Estimation Accuracy – PPG Database

Table I shows the average heart rate estimation error in bpm for each of the 12 subjects in the SP Cup database as well as the overall mean and standard deviation of error. We can see that the average error is < 2 bpm for most subjects. Also shown for comparison are the corresponding results from the JOSS, TROIKA and Robust EEMD works. Note that Table I shows the results for the truncated data, and results are presented for our proposed work as well as the Robust EEMD at both 25Hz and 125Hz sampling rate. The average errors are more or less similar for the different methods, with the ‘Robust EEMD’ marginally better, whereas the proposed method at 125Hz

shows the lowest standard deviation of error. The results for the un-truncated data are in Table II, and we can see that the error from the particle filter estimates are hardly affected despite the noisy initial periods that prohibited the use of the JOSS algorithm.

In Figure 6 below is shown the Bland-Altman plot for the particle filter estimates’ agreement with the ground truth at the full 125Hz sampling rate.




Figure 6 – Bland-Altman plot for particle filter agreement with ground truth

The limits of agreement (LOA) were defined following standard practice as $[\mu - 1.96\sigma, \mu + 1.96\sigma]$, where μ is the average difference and σ is the standard deviation, 2.35 bpm in this case. The LOA were $[-4.75, 4.45]$ bpm, and 95% of the difference values were within this confidence interval.

B. Heart Rate Estimation Accuracy – ECG Database and Fusion of ECG + PPG

Table III shows the estimation error when using the particle filter to estimate heart rate from the noisy ECG simulated as described in section V B. For comparison, we show the average estimation error for heart rates as computed by our implementation of the well-respected Pan-Tompkins algorithm, which was designed specifically to estimate heart rate from ECG signals [27]. Of course, the Pan-Tompkins algorithm was not designed for this intensity of motion artifacts, but we included it to show the extent of noisiness in the ECG which causes significant issues for an established algorithm. We can see how the particle filter also works well with this different modality with low error rates. In addition, also shown in the table are the results of fusion of this noisy ECG with the two noisy PPG channels and the accelerometer. We can see how the fusion almost always improves the accuracy, showing how the particle filter was able to effectively reward the consistent true observations across the different sources and make the best of the sensors available. The particle filter tracking over time for Subject 1 is also shown in Figure 7 for illustrative purposes. In this figure, ‘Findpeaks estimate’ refers to the heart rate estimate based solely on the CWT-based peak observation method on ECG, and it can be seen how it tends to overestimate as soon as the motion starts, whereas the particle filter continues to keep track even as the subject’s heart rate changes substantially during periods of motion activity.




Figure 7 – Heart rate estimation performance on a single subject

TABLE I. MEAN ABSOLUTE HEART RATE ESTIMATION ERROR (IN BPM) FOR THE VARIOUS ALGORITHMS ON THE TRUNCATED DATASETS

Subject #	1	2	3	4	5	6	7	8	9	10	11	12	Mean \pm SD
JOSS [14] (25Hz)	1.33	1.75	1.47	1.48	0.69	1.32	0.71	0.56	0.49	3.81	0.78	1.04	1.28 ± 2.61
TROIKA [13] (25Hz)	3.05	3.31	1.49	2.03	1.46	2.35	1.76	1.43	1.28	5.08	1.8	3.02	2.34 ± 2.86
Robust EEMD [15] (25Hz)	1.7	0.84	0.56	1.15	0.77	1.06	0.63	0.53	0.52	2.56	1.05	0.91	1.02 ± 1.79
Particle Filter (25Hz)(Our method)	2.21	1.71	1.11	1.71	1.1	1.72	1.11	1.29	1.12	3.5	1.68	1.57	1.65 ± 2.07
Robust EEMD [15] (125Hz)	1.83	0.85	0.63	1.21	0.65	1.03	0.7	0.5	0.47	2.83	1.14	0.9	1.06 ± 2.02
Particle Filter (125Hz) (Our method)	1.91	1.3	1.08	1.63	1.06	1.64	1.09	1.25	1.1	3.41	1.65	1.59	1.56 ± 1.73

TABLE II. MEAN ABSOLUTE HEART RATE ESTIMATION ERROR (IN BPM) FOR THE VARIOUS ALGORITHMS ON THE UN-TRUNCATED DATASETS

Subject #	1	2	3	4	5	6	7	8	9	10	11	12	Mean \pm SD
TROIKA [13] (25Hz)	3.05	3.49	1.49	2.03	1.46	2.35	1.76	1.42	1.28	5.73	1.79	3.02	2.41 ± 3.45
Robust EEMD [15] (25Hz)	1.64	0.81	0.57	1.44	0.77	1.06	0.63	0.47	0.52	2.94	1.05	0.91	1.07 ± 2.17
Particle Filter (25Hz) (Our method)	2.21	1.55	1.41	1.65	1.1	1.72	1.11	1.24	1.12	3.63	1.65	1.57	1.66 ± 2.17
TROIKA [13] (125Hz)	2.29	2.19	2	2.15	2.01	2.76	1.67	1.93	1.86	4.7	1.72	2.84	2.34 ± 0.82
Particle Filter (125Hz) (Our method)	1.91	1.46	1.39	1.61	1.06	1.64	1.09	1.25	1.1	3.58	1.73	1.59	1.62 ± 2.01

TABLE III. MEAN ABSOLUTE HEART RATE ESTIMATION ERROR (IN BPM) FOR THE ECG AND FUSION (ECG+PPG) PARTICLE FILTERS, AND PAN-TOMPKINS

Subject #	1	2	3	4	5	6	7	8	9	10	11	12	Mean \pm SD
ECG Particle Filter (Our method)	1.49	1.83	2.31	1.2	1.05	2.42	1.91	1.53	1.44	1.13	1.04	1.34	1.56 ± 2.02
ECG+PPG Particle Filter (Our method)	1.26	1.17	0.85	1.11	0.84	1.03	0.87	0.93	0.87	2.24	1.08	1.16	1.12 ± 1.32
Pan-Tompkins[27]	26.1	17.5	19.9	23.5	23.3	24.6	22.7	18.6	18.2	33.9	25.2	24.4	23.2 ± 20.02

C. Heart Rate Estimation – Experimental Data Collection

In Table IV we also present the results of heart rate error from the fusion particle filter on the dataset collected ourselves, which guarantees real simultaneous ECG and PPG affected by motion artifacts. This shows that the particle filter performance continues to be effective even in this scenario. Again, for comparison is shown the error rates when using the Pan-Tompkins algorithm on the noisy ECG. Note that for Subject 4 the Pan-Tompkins algorithm’s adaptive parameters completely went off track early on in the data record due to excessive noise, and did not recover estimates thereafter.

TABLE IV. MEAN ABSOLUTE ESTIMATION ERROR FOR FUSION PARTICLE FILTER AND PAN-TOMPKINS ON OUR EXPERIMENTAL DATASET

Subject #	Error for Particle Filter (bpm)	Error for Pan-Tompkins [27] (bpm)
1	1.55	13.73
2	1.63	19.57
3	1.25	11.46
4	1.12	N/A
5	1.47	10.4
Mean \pm SD	1.4 ± 1.55	13.79 ± 17.35

D. Fusion Contribution Analysis

In order to further illustrate how the fusion of modalities works, we take a closer look at the performance on Subject 10 from the database. As can be seen in Tables I and II, estimation performance on this subject is noticeably worse, for our algorithm as well as those of other previous works. This suggests that the PPG signals themselves were relatively more unreliable for this subject. However, we see that in Table III when using the noisy ECG the performance is much better; so we can assume in this instance that the ECG is a more reliable signal at least for certain segments of the data.

Figure 8 shows the relative contribution of each modality – ECG and the two PPG sensors – over time for Subject 10, computed as described in equations (26) and (27). In this

figure, we plot only a subset of the time windows, spanning about 1 minute. Moreover, overlaid in red is the particle filter heart rate estimation error for each of those windows. The error rises to almost 20 beats per minute around window 10, but soon after this the contribution of the ECG to the overall estimate increases. It is clear that the particle filter fusion rewards the more consistent observations from the ECG, and correspondingly the overall error drops sharply. We see a similar trend on a smaller scale around time window 40, where the error is relatively high until the ECG contributions become higher and the overall estimation performance becomes better. In future work, we aim to implement techniques that can recognize these trends of quality of observations and explicitly re-weight individual modalities in the fusion formulation.




Figure 8 – Relative contribution of ECG and PPG modalities to overall fusion particle filter estimate over time for Subject #10

E. Discussion of Estimation Performance

The estimation errors are low, but in order to provide further context, we have compared the results to those of recent state-of-the-art works on heart rate estimation in the presence of motion artifacts. The estimation error levels are comparable to the most recent related works in the area. We note that the other related works were specifically developed and optimized for the objective of heart rate monitoring using PPG signals with several heuristics; for instance, TROIKA and JOSS use heuristics such as a rigid artificial bound on the variability of reported heart rate estimates from one window to the next, thresholds for what constitutes a big enough peak in the PPG frequency spectrum, and polynomial curve fitting based on previous heart rate estimates to predict the next estimate when the tracking does not return a satisfactory result. Similarly, the ‘Robust EEMD’ work, in addition to

using EEMD and an adaptive filter, has an arbitrary ‘absolute criterion’ to designate a ‘reliable peak’ in the PPG spectrum for heart rate and thresholds for what constitutes a strong enough peak in the PPG spectrum. The algorithm also deletes or removes segments of the signal from consideration if the corresponding accelerometer magnitude is too high. Moreover, with the EEMD approach the user is required to manually detect in a training phase which of the several intrinsic mode functions has the pertinent heart rate frequency information, and this also changes with sampling rate. It is therefore notable that the relatively more generalized particle filter framework introduced here with minimal heuristics or rule-based steps, no requirement for clean start, no deletion of data, which can work with other signal modalities as shown with ECG, and can also be applied to other physiological signal estimation problems, exhibits comparable performance to contemporary works that were purpose-built for the heart rate estimation problem on PPG signals. Moreover, as will be noted in the next section, this comparable estimation performance is achieved with an algorithm that is far more computationally efficient compared to these works.

F. Computation Time

In this work, due to the formulation with the heart rate state, we mitigated the computational load by tracking only one state dimension with just 300 particles. Indeed, the contemporary works we can compare this to are significantly more computationally intensive. The authors of the ‘Robust EEMD’ work [15] note that the TROIKA algorithm takes about 17 minutes and 30 seconds on average to complete heart rate estimation on a single subject at a sampling rate of 125Hz; whereas the Robust EEMD algorithm itself takes about 55 seconds per subject. Similarly, at a sampling rate of 25Hz, the JOSS algorithm takes about 25s on average per subject and the corresponding Robust EEMD algorithm takes about 16s. When we measured the execution time of our particle filter implementation on MATLAB, the average time per subject was only about 1.04 seconds for the 25Hz sampling rate, and 1.18 seconds for the 125Hz sampling rate. It must be noted that the execution times reported above for the related works were gathered from a work that used MATLAB 2013a, whereas we use MATLAB 2017a. However, this alone cannot account for the highly significant difference in computation time. Furthermore, the machine used to extract these results has similar specifications to the one used to report the results for the related works [15]. In particular, we used a Windows 10 64-bit PC with an Intel i7-6700 processor at 2.60 GHz and 16GB of RAM.

We also analyzed the trade-off between the accuracy and computational cost as a function of the number of particles. Figure 9 shows a comparison of the error rates and computation time per minute of data for our particle filter as the number of particles is varied for a single subject. As a reminder, we used $N = 300$ particles in our work. While the estimation performance does improve as we increase the number of particles, as expected, it is likely that the higher values of N would make it impractical to compute these estimates in real-time, especially on wearable sensors. On such systems, one can easily adjust the number of particles subject to the availability of computational resources.




Figure 9 – Changes in average estimation error and computation time per minute of data on a single subject as the number of particles is changed

G. Limitations

We did not test on patients with heart rate variability or other cardiac conditions; this will likely require some tuning of the parameters, but this would be applicable to other contemporary signal processing techniques as well. Testing on subjects with abnormal cardiac activity will be left for future work. We also note that precise computational benchmarking is not the primary goal of this work; the previous section was only meant to provide a rough guide indicating a definite computational advantage over contemporary related works in the area. Deployment of the algorithm on a system is out of the scope of this work; however we submit that the design of such a system is eminently feasible, especially if we leverage cloud computing resources or other techniques to circumvent the computational constraints on typical wearable systems.

VII. CONCLUSION

In this work, we have introduced a generalized particle filter framework that can be used to track heart rate and proved the feasibility of the technique on real world PPG and ECG signals affected by motion artifacts. Furthermore, we showed how the particle filter can be used to successfully improve estimation accuracy by combining information from multiple modalities simultaneously measuring the same target phenomenon or even the noise associated with the target. This will prove useful in the context of the upcoming IoT ecosystem where there are multiple wearable and environmental sensors continuously monitoring the physiological status of the user.

ACKNOWLEDGMENT

This work was supported in part by the National Science Foundation, under grants CNS-1150079 and EEC-1648451, and by TerraSwarm, one of six centers of STARnet, a Semiconductor Research Corporation program sponsored by MARCO and DARPA. Any opinions, findings, conclusions, or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the funding organizations.

REFERENCES

- [1] G. E. Prinsloo, H. L. Rauch, and W. E. Derman, “A brief review and clinical application of heart rate variability biofeedback in sports, exercise, and rehabilitation medicine,” *The Physician and Sportsmedicine*, vol. 42, no. 2, pp. 88–99, 2014.
- [2] F. Lin, C. Song, Y. Zhuang, W. Xu, C. Li, and K. Ren, “Cardiac Scan: A non-contact and continuous heart-based authentication system,” in *2017 ACM International Conference on Mobile Computing and Networking (MobiCom)*, October 2017.
- [3] S. Ha, C. Kim, Y. M. Chi, A. Akinin, C. Maier, A. Ueno, and G. Cauwenberghs, “Integrated circuits and electrode interfaces for

- noninvasive physiological monitoring," *IEEE Transactions on Biomedical Engineering*, vol. 61, pp. 1522–1537, May 2014.
- [4] S. Hwang, J. Seo, H. Jebelli, and S. Lee, "Feasibility analysis of heart rate monitoring of construction workers using a photoplethysmography (PPG) sensor embedded in a wristband-type activity tracker," *Automation in Construction*, vol. 71, no. Part 2, pp. 372–381, 2016.
 - [5] O. Cappe, S. Godsill, and E. Moulines, "An overview of existing methods and recent advances in sequential Monte Carlo," *Proceedings of the IEEE*, vol. 95, pp. 899–924, May 2007.
 - [6] A. Hennig and A. Patzak, "Continuous blood pressure measurement using pulse transit time," *Somnologie - Schlaforschung und Schlafmedizin*, vol. 17, no. 2, pp. 104–110, 2013.
 - [7] C. Yang and N. Tavassolian, "Motion noise cancellation in seismocardiogram of ambulant subjects with dual sensors," in *2016 38th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, pp. 5881–5884, Aug 2016.
 - [8] D. Jarchi and A. J. Casson, "Estimation of heart rate from foot worn photoplethysmography sensors during fast bike exercise," in *2016 38th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, pp. 3155–2158, Aug 2016.
 - [9] N. V. Helleputte, M. Konijnenburg, J. Pettine, D. W. Jee, H. Kim, A. Morgado, R. V. Wegberg, T. Torfs, R. Mohan, A. Breeschoten, H. de Groot, C. V. Hoof, and R. F. Yazicioglu, "A 345 μ W multi-sensor biomedical SoC with bio-impedance, 3-channel ECG, motion artifact reduction, and integrated DSP," *IEEE Journal of Solid-State Circuits*, vol. 50, pp. 230–244, Jan 2015.
 - [10] A. Galli, G. Frigo, C. Narduzzi, and G. Giorgi, "Robust estimation and tracking of heart rate by PPG signal analysis," in *2017 IEEE International Instrumentation and Measurement Technology Conference (I2MTC)*, pp. 1–6, May 2017.
 - [11] K. Sweeney, T. Ward, and S. McLoone, "Artifact removal in physiological signals - practices and possibilities," *Information Technology in Biomedicine, IEEE Transactions on*, vol. 16, pp. 488–500, May 2012.
 - [12] S. S. Bisht and M. P. Singh, "An adaptive unscented Kalman filter for tracking sudden stiffness changes," *Mechanical Systems and Signal Processing*, vol. 49, no. 1–2, pp. 181–195, 2014.
 - [13] Z. Zhang, Z. Pi, and B. Liu, "TROIKA: A general framework for heart rate monitoring using wrist-type photoplethysmographic signals during intensive physical exercise," *Biomedical Engineering, IEEE Transactions on*, vol. 62, pp. 522–531, Feb 2015.
 - [14] Z. Zhang, "Photoplethysmography-based heart rate monitoring in physical activities via joint sparse spectrum reconstruction," *IEEE Transactions on Biomedical Engineering*, vol. 62, pp. 1902–1910, Aug 2015.
 - [15] E. Khan, F. A. Hossain, S. Z. Uddin, S. K. Alam, and M. K. Hasan, "A robust heart rate monitoring scheme using photoplethysmographic signals corrupted by intense motion artifacts," *IEEE Transactions on Biomedical Engineering*, vol. 63, pp. 550–562, March 2016.
 - [16] S. Kim, L. A. Holmstrom, and J. McNames, "Tracking of rhythmical biomedical signals using the maximum a posteriori adaptive marginalized particle filter," *British Journal of Health Informatics and Monitoring*, vol. 2, no. 1, 2015.
 - [17] C. Lin, M. Bugallo, C. Mailhes, and J.-Y. Tournet, "ECG denoising using a dynamical model and a marginalized particle filter," in *Signals, Systems and Computers (ASIOMAR), 2011 Conference Record of the Forty Fifth Asilomar Conference on*, pp. 1679–1683, Nov 2011.
 - [18] G.-J. Li, X. na Zhou, S. ting Zhang, and N.-Q. Liu, "ECG characteristic points detection using general regression neural network-based particle filters," in *Bioelectronics and Bioinformatics (ISBB), 2011 International Symposium on*, pp. 155–158, Nov 2011.
 - [19] G. Li, X. Zeng, J. Lin, and X. Zhou, "Genetic particle filtering for denoising of ECG corrupted by muscle artifacts," in *Natural Computation (ICNC), 2012 Eighth International Conference on*, pp. 562–565, May 2012.
 - [20] S. Edla, N. Kovvali, and A. Papandreou-Suppappola, "Sequential Markov chain Monte Carlo filter with simultaneous model selection for electrocardiogram signal modeling," in *Engineering in Medicine and Biology Society (EMBC), 2012 Annual International Conference of the IEEE*, pp. 4291–4294, Aug 2012.
 - [21] I. R. Legarreta, P. S. Addison, M. J. Reed, N. Grubb, G. R. Clegg, C. E. Robertson, and J. N. Watson, "Continuous wavelet transform modulus maxima analysis of the electrocardiogram: beat characterisation and beat-to-beat measurement," *International Journal of Wavelets, Multiresolution and Information Processing*, vol. 03, no. 01, pp. 19–42, 2005.
 - [22] V. Nathan, I. Akkaya, and R. Jafari, "A particle filter framework for the estimation of heart rate from ECG signals corrupted by motion artifacts," in *Engineering in Medicine and Biology Society (EMBC), 2015 37th Annual International Conference of the IEEE*, pp. 6560–6565, Aug 2015.
 - [23] F. Caron, M. Davy, E. Duflos, and P. Vanheeghe, "Particle filtering for multisensor data fusion with switching observation models: Application to land vehicle positioning," *IEEE Transactions on Signal Processing*, vol. 55, pp. 2703–2719, June 2007.
 - [24] A. L. Goldberger, L. A. N. Amaral, L. Glass, J. M. Hausdorff, P. C. Ivanov, R. G. Mark, J. E. Mietus, G. B. Moody, C.-K. Peng, and H. E. Stanley, "PhysioBank, PhysioToolkit, and PhysioNet: Components of a new research resource for complex physiologic signals," *Circulation*, vol. 101, no. 23, pp. e215–e220, 2000 (June 13).
 - [25] G. B. Moody, W. Muldrow, and R. G. Mark, "A noise stress test for arrhythmia detectors," *Computers in Cardiology*, vol. 11, pp. 381–384, 1984.
 - [26] S. S. Thomas, V. Nathan, C. Zong, K. Soundarapandian, X. Shi, and R. Jafari, "Biowatch: A noninvasive wrist-based blood pressure monitor that incorporates training techniques for posture and subject variability," *IEEE Journal of Biomedical and Health Informatics*, vol. 20, pp. 1291–1300, Sept 2016.
 - [27] J. Pan and W. J. Tompkins, "A real-time QRS detection algorithm," *IEEE Transactions on Biomedical Engineering*, vol. BME-32, pp. 230–236, March 1985.



Viswam Nathan (M'14) received his B.S. and M.S. degrees in computer engineering from UT-Dallas in 2012 and 2015 respectively. He is currently working toward his Ph.D. in computer engineering at Texas A&M University. His research interests include design and development of wearable and reconfigurable health monitoring devices, and associated signal processing techniques.



Roozbeh Jafari (SM'12) is an associate professor in Biomedical Engineering, Computer Science and Engineering and Electrical and Computer Engineering at Texas A&M University. He received his Ph.D. in Computer Science from UCLA and completed a postdoctoral fellowship at UC-Berkeley. His research interest lies in the area of wearable computer design and signal processing. His research has been funded by the NSF, NIH, DoD (TATRC), AFRL, AFOSR, DARPA, SRC and industry (Texas Instruments, Tektronix, Samsung & Telecom Italia). He has published over 100 papers in refereed journals and conferences. He has served as the general chair and technical program committee chair for several flagship conferences in the area of Wearable Computers. He is the recipient of the NSF CAREER award in 2012, IEEE Real-Time & Embedded Technology & Applications Symposium (RTAS) best paper award in 2011 and Andrew P. Sage best transactions paper award from IEEE Systems, Man and Cybernetics Society in 2014. He is an associate editor for the IEEE Transactions on Biomedical Circuits and Systems, IEEE Sensors Journal, IEEE Internet of Things Journal and IEEE Journal of Biomedical and Health Informatics. He serves on scientific panels for funding agencies frequently and is presently serving as a standing member of the NIH Biomedical Computing and Health Informatics study section.