A Particle Filter Framework for the Estimation of Heart Rate from ECG Signals Corrupted by Motion Artifacts

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

***Abstract*— In this work, we describe a methodology to probabilistically estimate the R-peak locations of an electrocardiogram (ECG) signal using a particle filter. This is useful for heart rate estimation, which is an important metric for medical diagnostics. Some scenarios require constant in-home monitoring using a wearable device. This poses a particularly challenging environment for heart rate detection, due to the susceptibility of ECG signals to motion artifacts. In this work, we show how the particle filter can effectively track the true R-peak locations amidst the motion artifacts, given appropriate heart rate and R-peak observation models. A particle filter based framework has several advantages due to its freedom from strict assumptions on signal and noise models, as well as its ability to simultaneously track multiple possible heart rate hypotheses. Moreover, the proposed framework is not exclusive to ECG signals and could easily be leveraged for tracking other physiological parameters. We describe the implementation of the particle filter and validate our approach on real ECG data affected by motion artifacts from the MIT-BIH noise stress test database. The average heart rate estimation error is about 5 beats per minute for signal streams contaminated with noisy segments with SNR as low as -6 dB.**

I. Introduction

Of all the physiological parameters associated with the human body, those related to the heart consistently draw the most attention. One of the most widely used indicators of heart health is the heart rate obtained through the electrocardiogram (ECG). Not only is this a standard vital sign monitored in hospitals, it is also useful for recovering patients to self-monitor in-home. In addition, the heart rate also provides valuable feedback to health enthusiasts, especially during exercise.

The scenario of in-home, around the clock monitoring is of particular interest, as demonstrated by the recent proliferation of wearable monitoring devices on the market. Heart rate variability (HRV) for instance, is an important indicator of heart health and it is a trend that needs to be monitored throughout the day; for example it has been shown that reduced HRV could predict the occurrence of myocardial ischemia [1].

Monitoring in day-to-day life in a home setting or during exercise presents significant obstacles. Firstly, long-term in-home monitoring would require a comfortable, wearable solution which would preclude the use of more stable and reliable electrode contacts such as gel-based adhesive patches. Secondly, there will likely be motion artifacts affecting the stream of ECG due to the various activities performed by the user during the course of the day. Both of these factors result in a noisier ECG stream and estimating the heart rate from this is not so straightforward.

As an example, Figure 1 shows two versions of the same ECG segment, one clean and the other corrupted by motion artifacts with 0 dB SNR. In both cases, the true ECG R-peaks are marked in blue. The bottom figure annotates all the peaks detected by a wavelet-based peak-detection algorithm, including false positives that are marked in black.

There have been many proposed approaches in the literature to obtain an accurate heart rate estimate from a noisy ECG signal. Several works have been based on the use of an adaptive filter, but such techniques always depend on the availability of an external reference signal, such as accelerometer data [2] or electrode tissue impedance [3], which in turn increases the complexity and cost of the hardware. 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 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 [4]. The extended Kalman filter was introduced to circumvent the disadvantage of the linearity assumption [5], but just like the regular Kalman filter it still suffers from the fact that only unimodal Gaussian distributions can be tracked [4]. In other words, only one state can be tracked at a time and if the estimate diverges from the true state, it may continue to diverge beyond recovery.

Blind source separation methods such as independent component analysis (ICA) used for motion artifact reduction [6] depend on having multiple streams of data, and again, this brings extra hardware cost; not to mention potentially compromising the comfort of the user in having to wear multiple sensors throughout the day.

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. It is robust and has the potential to recover from incorrect estimates since it always keeps track of multiple possibilities. It is generalizable and can relatively easily be adapted to handle a variety of signal and noise models. It is also straightforward to tune certain parameters of the particle filter to trade-off

V. Nathan and R. Jafari are with the Electrical Engineering Department, University of Texas at Dallas, Richardson, TX 75080 USA (e-mail: {viswamnathan,rjafari}@utdallas.edu).

Ilge Akkaya is with the EECS Department, University of California at Berkeley, Berkeley, CA 94720 USA (email: ilgea@eecs.berkeley.edu)

between computational complexity, accuracy and robustness depending on the application scenario.

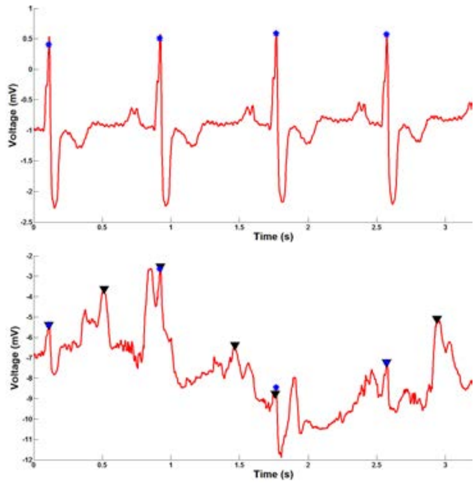

Figure 1 - Clean ECG with R-peaks marked (top) and Motion Artifact affected ECG (bottom)

The particle filter has been previously employed in other similar applications. However these usually involve a complex dynamical model for the ECG, which involves several state dimensions which in turn increases the computational cost and complexity [7-9]. 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 [10]. The target application scenario for these is detecting characteristic points in clean or white-noise affected ECG, which is quite different from our focus here.

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 [11]. 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 ECG signals from different subjects [12].

The primary contribution of this work is the introduction of an intuitive and lightweight particle filter framework that processes noisy ECG by probabilistically determining if a given peak in the signal stream corresponds to an R-peak or an artifact, followed by validation of the approach on real motion artifact affected ECG data. Moreover, the framework itself is generic enough such that it could be applicable to a wide range of users and scenarios as well as the tracking of similar physiological signals other than ECG, such as the photoplethysmogram (PPG).

II. Theory

A. Particle Filter

The particle filter employs a recursive technique to estimate the current state of the system by iteratively propagating and weighting a set of particles such that they converge to represent the posterior probability density of the state. In the current work, the state being estimated is the current heart rate corresponding to the ECG signal. Each particle represents a possible heart rate and the weight distribution of particles indicates the most probable candidate for the true heart rate of the signal; particles with higher weight have a higher probability of being the true state of the system. This distribution is guided by a weight update step, based on information from an observation mechanism, *i.e.* some method to observe the current true heart rate via noisy measurements, and provide a measure of the correctness of that observation. The observation mechanism will be described in detail in Section III B.

B. Features and Limitations

The particle filter requires no assumptions about the probability distribution of the state space, the observation model, or the nature of the noise affecting these. This allows the filter to track multiple hypotheses within an arbitrarily distributed probability space. This feature is especially useful for ECG data because very often, motion of the sensor causes spike-like artifacts which could be mistaken for the R-peak of ECG. Consequently, the posterior heart rate probability distribution, based on whether or not false positives are included, actually becomes a multi-modal distribution, rather than a single Gaussian distribution around the true value. The particle filter can keep track of these multiple possible heart rates until it converges to the true estimate.

One of the primary disadvantages of using a particle filter is the computational effort; for several applications, especially those that rely on low-power devices with limited computational resources, handling a large number of particles may be unfeasible. For our application, however, we are only looking at tracking one dimension of data and we can achieve a low error rate with a relatively small number of particles compared to other implementations.

III. Methods

A. ECG Model

Figure 2 shows two beats of a clean ECG stream with the R-peak and R-R interval marked. In this work, the particle filter probabilistically estimates the locations of R-peaks by leveraging the properties of consecutive R-R intervals; namely the fact that changes in the heart rate of the human heart are relatively smooth, stable and slow-changing.

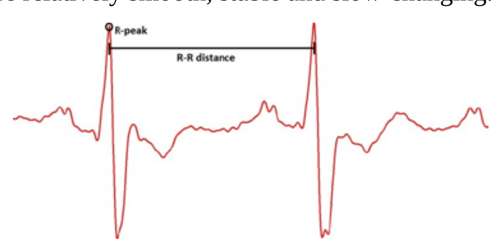

Figure 2 - ECG signal showing R-peak and R-R interval

The particle filter estimates the heart rate for each fixed length moving window of the continuous ECG stream. If $ecg(t)$ represents the time domain ECG signal with duration of T seconds, we can define the window parameters as follows:

$$numW = (T - T_{duration})/T_{step} \quad (1)$$

$$t_1 = (i-1)T_{step} \quad (2)$$

$$t_2 = (i-1)T_{step} + T_{duration} \quad (3)$$

$$W_i = ecg(t_i),\ t_1 \leq t_i \leq t_2 \quad (4)$$

$$\forall i \in (1, numW),$$

where $numW$ is the total number of windows, $T_{duration}$ is the window time duration (in seconds), T_{step} is the window step size (in seconds) and W_i is the i^{th} window of ECG

Keeping in mind the calculation complexity, desired resolution of heart rate and the sufficiency of information in a window, we heuristically fixed the window time period and the window step size to be 4 seconds each.

We assume in this model that the true heart rate does not vary by a few beats per minute (bpm) between two consecutive non-overlapping windows, which is a heuristic assumption. Hence, in this initial study, we do not yet consider extreme situations where higher rates of change are observed.

B. Observation Mechanism

The observation mechanism is used in particle filters to give some insight into the current true state of the system and it dictates the weight or reliability of each of the observations.

In our work this observation mechanism is used to identify all possible candidates for an R-peak in one window of ECG. This is based on the MATLAB 'findpeaks' function as follows:

$$findpeaks(W_i, minAmp, minTime)$$

$$= \{P_1, P_2, \dots, P_k\} = P_{W_i}$$

$$\forall i \in (1, numW)$$

where $findpeaks(W_i, minAmp, minTime)$ finds the time of occurrence of all peaks in W_i that have amplitude at least minAmp and such that no two peaks are within minTime distance, and P_{W_i} is the set of peak locations in time, $\{P_1, P_2, \dots, P_k\}$, returned by the findpeaks function.

We then calculate the heart rate (in bpm) corresponding to each and every possible combination of these peaks in that window. The number of combinations for each window is given by:

$$NC_i = \sum_{j=3}^{k_i} \binom{k_i}{j} \quad (5)$$

$$PC_i \triangleq The\ set\ of\ NC_i\ combinations\ of\ elements\ of P_{W_i}$$

$$\forall i \in (1, numW)$$

where NC_i is the total number of combinations of peak locations in window W_i, and k_i is the number of peak locations in the set P_{W_i}. The set of combinations is restricted to those involving at least 3 peaks due to the nature of the observation weighting procedure, which is described at the end of this section.

The observations are the heart rates estimates for each of these peak combinations, and are calculated as follows:

$$diff_j = {PC_i}^n(j+1) - {PC_i}^n(j) \quad (6)$$

$$avgInterval = \frac{\sum_{j=1}^{m-1} diff_j}{m-1} \quad (7)$$

$$HRObserve_i^{\ n} = \left(\frac{1}{avgInterval}\right) \times 60 \quad (8)$$

$$\forall i \in (1, numW), \forall n \in (1, NC_i), \forall j \in (1, m-1)$$

where ${PC_i}^n(j)$ is the j^{th} peak location in the n^{th} element of PC_i, m is the number of peak locations in ${PC_i}^n$ and $HRObserve_i^{\ n}$ is the heart rate observation in bpm corresponding to the n^{th} combination of peak locations in window W_i

Each of these heart rate candidates is then assigned a weight, which reflects the confidence that the set of peaks corresponding to that heart rate consists of true R-peaks only. In our work, we wanted to leverage the fact that the human heart rate is relatively steady and slow-changing, whereas peaks caused by motion artifacts are likely to be irregularly spaced in time. Keeping these in mind, the weight measure for a given set of peaks was calculated as follows:

$$\sigma_i^{\ n} = stddev(\forall j \in (1, m-1), diff_j) \quad (9)$$

$$obsWeight_i^{\ n} = 1/\sigma_i^{\ n} \quad (10)$$

$$\forall i \in (1, numW), \forall n \in (1, NC_i)$$

where $\sigma_i^{\ n}$ is the standard deviation of the set of all elements of $diff_j$ for the n^{th} combination of peak locations of the i^{th} window and $obsWeight_i^{\ n}$ is the weight assigned to the n^{th} combination of peak locations in window W_i

Since this mechanism relies on standard deviation and seeing a trend in the heart rate variability, we decided that only combinations of at least 3 peaks would be considered in each window. Note that since the window size is 4 seconds, we expect at least 3 true R-peaks to be present even for very low heart rates. The proposed approach hinges on the assumption that at least one of the combinations constitutes the set of true R-peaks.

This 'findpeaks' function is a little too naïve on its own for noisy signals and so we leveraged the continuous wavelet transform (CWT) to reduce the number of false positives. We applied the CWT on the noisy ECG signal, and then used the 'findpeaks' function on the resulting transformed signal (with higher thresholds) to find the locations of the peaks. According to the guidance of a previous work, a Mexican Hat wavelet with a center frequency of 0.25 Hz and a scale of 5.29 was used to mimic the shape of an R-peak [13]. Each of the windows W_i referred to in this work are the windows of the wavelet transformed ECG signal.

It must be noted that although this process eliminates a lot of the false positives, several of them still remain and the particle filter is still necessary to track the true heart rate. It will be discussed further in Section IV that the heart rate estimation performance suffers if we rely solely on threshold based peak detection even after the wavelet transformation.

C. Particle Filtering

Estimating the true heart rate given a noisy ECG signal can be formulated as a *state estimation* problem, where the state space representation of the system is given by:

$$\mathcal{X}_t \sim \pi_x(\mathcal{X}_t) \quad \text{(prior distribution)}$$
$$Z_t|\mathcal{X}_t \sim g(\mathcal{X}_t) \quad \text{(measurement model)}$$
$$\mathcal{X}_{t+1}|\mathcal{X}_t \sim f(\mathcal{X}_t) \quad \text{(state dynamics)}$$

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 probability distribution of the system based on prior knowledge, Z_t denotes a set of discrete observations, *i.e.*, the set of heart rate observations as described in the previous section, $g(\cdot)$ is a function representing the observations conditioned on the true heart rate, and $f(\cdot)$ is the state dynamics model that characterizes the heart rate dynamics as a function of the moving window indices of the ECG signal 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 how likely the particle is to being 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:

$$p(\mathcal{X}_t|Z_t) = \sum_{p=1}^{N_p} W_{X_t^p} \delta(\mathcal{X}_t - X_t^p) \qquad (11)$$

where X_t^p is the p^{th} particle at window *t*, $W_{X_t^p}$ denotes the weight of particle 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.

The state-space representation of the heart-rate estimation problem given a noisy ECG signal can be formulated as follows:

$$\mathcal{X}_t \sim Uniform(HR_{min}, HR_{max}) \quad \text{(prior distribution)}$$
$$Z_t|\mathcal{X}_t \sim g(\mathcal{X}_t) \quad \text{(measurement model)}$$
$$\mathcal{X}_{t+1}|\mathcal{X}_t \sim N(\mathcal{X}_t, \sigma_x) \quad \text{(state dynamics)}$$

where HR_{min} and HR_{max} denote the assumed lower and upper limits of the heart rate. The prior distribution is considered to be a uniform distribution between these two limits. Here, $N(\mathcal{X}_t, \sigma_x)$ denotes a Gaussian distribution with mean equal to the true heart rate in window *t*, and standard deviation σ_x reflecting the maximum expected change in heart rate from one window to the next. Thus, the state dynamics of the heart rate is formulated as a Gaussian random walk between windows and σ_x is set to be 3 bpm in accordance with the ECG model described earlier in Section III A.

The measurement model $g(\mathcal{X}_t)$, which relates the observations to the true heart rate, can be modeled as follows:

$$p(Z_t|\mathcal{X}_t) = \sum_{n=1}^{NC_i} p(Z_t^n|\mathcal{X}_t)$$
$$= \sum_{n=1}^{NC_i} obsWeight_t{}^n N(\mathcal{X}_t, \sigma_z) \qquad (12)$$

where Z_t^n refers to $HRObserve_t{}^n$, the n^{th} heart rate estimate in window *t* from (8), $obsWeight_t{}^n$ is the weight associated with the observation, which can be thought of as a likelihood measure of the observation belonging to the true heart rate, and $N(\mathcal{X}_t, \sigma_z)$ denotes a Gaussian distribution with mean equal to the heart rate in window *t*, and standard deviation σ_z reflecting the maximum tolerable deviation between the true heart rate and the observation.

The particles are initially distributed uniformly between HR_{min} and HR_{max}, defined to be 30 and 220 bpm respectively for this work. The total number of particles, N_p, is set to be 100 and the individual particle weights are all set to be $\frac{1}{N_p}$.

When the first window of ECG is processed, each particle is assigned a weight to reflect the probability that it represents the true heart rate. This is done as follows:

$$W_{X_i^p} \propto p(Z_i|X_i^p) = \sum_{n=1}^{NC_i} p(Z_i^n|X_i^p)$$
$$= \sum_{n=1}^{NC_i} obsWeight_i{}^n \times N(Z_i^n, X_i^p, \sigma_z) \qquad (13)$$
$$WNorm_{X_i^p} = W_{X_i^p} / \sum_{r=1}^{N_p} W_{X_i^r} \qquad (14)$$
$$\forall p \in (1, N_p)$$

where X_i^p is the p^{th} particle of the i[th] window, $W_{X_i^p}$ is the weight of particle X_i^p, $N(Z_i^n, X_i^p, \sigma_z)$ is the value of a Gaussian distribution with mean X_i^p and standard deviation σ_z evaluated at Z_i^n, and $WNorm_{X_i^p}$ is the normalized weight of the p^{th} particle of the i[th] window. σ_z is set to be 2 bpm.

The heart rate corresponding to the particle with the maximum weight is then returned by the particle filter as its heart rate estimate for that window:

$$HREst_i = X_i^m \Big|^{argmax}_{m} WNorm_{X_i^m} \qquad (15)$$

where $HREst_i$ is the heart rate estimate of the particle filter in bpm for window W_i and m is the index of the particle with the maximum weight.

All the particles are then re-sampled such that, while maintaining the same total number of particles, there are now more particles around heart rates that had higher weight particles in the previous step. This is done in accordance with the well-known sampling importance resampling (SIR) procedure of particle filters to avoid the problem of particle degeneracy [14]. as follows:

$$XSum_i^p = \sum_{r=1}^{N_p} WNorm_{X_i^r} \qquad (16)$$
$$u = {}^{argmin}_{a} | rand(0,1) \leq XSum_i^a \qquad (17)$$
$$XUpdate_i^p = X_i^u \qquad (18)$$
$$\forall p \in (1, N_p)$$

where $XUpdate_i^p$ is the updated state of the p^{th} particle of the i[th] window after resampling and $rand(0,1)$ is a randomly generated number from a uniform distribution between 0 and

1. In the new distribution of particles, the weights are all reset to be equal.

While transitioning to the next window, all particles are propagated according to the state dynamics model as follows:

$$X_{i+1}{}^{p} \sim p(X_{i+1}^{p} | X_{i}^{p})$$
$$\sim N(XUpdate_{i}^{p}, \sigma_x)$$
$$= XUpdate_{p}{}^{i} + (\sigma_x \cdot randn) \qquad (19)$$
$$\forall p \in (1, N_p)$$

where $randn$ is a randomly generated number from the standard normal distribution and σ_x is the heart rate deviation limit defined earlier.

This is done to ensure that the particles reflect the changing state of the true heart rate and do not stagnate on values converged on previously. The entire process is repeated with the next window of ECG on this new distribution of particles.

IV. Experiments and Results

A. ECG Database

The ECG data is taken from the MIT-BIH Arrhythmia Database and the corresponding noise is taken from the MIT-BIH Noise Stress Test Database [15, 16]. The noise data includes electrode motion artifacts and the database itself provides techniques to reliably add known amounts of this noise to the corresponding clean ECG data. In order to ensure that the various SNR levels are meaningful with respect to heart rate detection, the database defines signal power in terms of the QRS amplitude and the noise power is measured in a frequency weighted manner.

For this work, we considered only those subjects whose ECG streams allowed for relatively simple R-peak detection for clean data. In other words, subjects exhibiting frequent ventricular ectopies or highly elevated T-segments were considered out of the scope of this initial study. With this criteria, we picked 7 subjects from the database; namely Subjects 103, 112, 115, 117, 122, 123 and 230.

In this work, we evaluate the proposed particle filter on ECG signals with SNR levels ranging from 6dB to -6dB. The criterion for evaluation of performance is the error with respect to the true heart rate which was obtained using a simple peak detection algorithm on the clean ECG signal.

For the 'findpeaks' function, the parameter minAmp was set to be 40% of the maximum amplitude of the corresponding clean signal, and minTime was set to be about 270ms, corresponding to a heart rate of about 220 bpm which we consider to be the maximum in this target application.

B. Heart Rate Estimation Results

The signals were first high-pass filtered in MATLAB above 0.5Hz to remove the baseline wander. Each subject's data consisted of a continuous ECG stream of about 30 minutes. The first 5 minutes was clean, followed by alternating 2-minute segments of continuously noisy data (at the defined SNR level) and 2-minute segments of continuously clean data. This entire stream was processed by the particle filter, with the heart rate reported for each 4-second window.

Table I shows the overall heart rate estimation error of the particle filter, averaged over all time for each subject, as well as averaged over all 7 subjects, for the various SNR levels. Since there is some randomness associated with the distribution of particles in each step, the results are slightly different for every run of the filter on the same data. Therefore, the particle filter error numbers below are the result of averaging after 5 such runs on each dataset. Also listed in the table is the error from using the 'CWT Findpeaks' estimate, which entails calculating the heart rate from the peaks returned by the findpeaks function after the wavelet transform. This gives us a measure of the system performance if it relied only on the time-series observation mechanism and not the particle filter that follows.

TABLE I. Mean Absolute Heart Rate Estimation Error

SNR	Particle Filter error (bpm)	CWT Findpeaks error (bpm)
6dB	1.402	12.475
3dB	2.169	17.613
0dB	3.461	21.798
-3dB	4.442	24.742
-6dB	5.044	27.572

C. Discussion of Results

As can be seen from the results of Table I, the particle filter exhibits competitive performance even for SNR as low as -6dB and it compares favorably to relying on wavelet transform based peak detection.

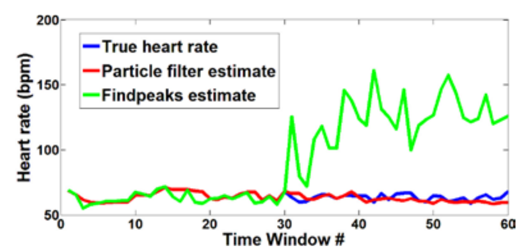

Figure 3 - Particle filter tracking compared to CWT Findpeaks method

Figure 3 shows the tracking performance of the particle filter over 4 minutes for Subject 115's ECG stream. The first 2 minutes are clean followed by 2 minutes of noisy data with 0dB SNR. We can see that while the CWT findpeaks method tracks perfectly during the clean sections, it drastically overestimates due to its inability to distinguish the false positives caused by motion artifacts in the latter half of the data. The particle filter estimate on the other hand, remains relatively stable and accurate in its estimate.

While these results are promising, there are some limitations with the current approach. Firstly, while the standard deviation based weighting scheme brings the advantage of an intuitive mechanism for this initial study, it may be too simplistic for extremely noisy sections. Secondly, the method of examining every combination of peaks exhaustively within a window has no theoretical upper bound and could bring about significant computational wastage due to its naivety. We intend to address both of these issues as well as continue to explore the potential of the particle filter for physiological signal monitoring applications, especially the implementation in low power embedded systems, in future works.

V. Conclusion

In this work we present a particle filter framework for obtaining the heart rate from a noisy, motion artifact affected ECG signal. Unlike other techniques, a statistical particle filter based approach does not make any assumptions about the type of noise nor of the distribution of the state being estimated. Testing on real motion artifact affected ECG data is shown to have average error of only around 5bpm even during conditions with SNR as low as -6dB. This represents a very feasible solution for continuous and pervasive monitoring of heart rate. Moreover, due to the generic nature of the framework it is our hope that in future this technique can be leveraged for other physiological signals and applications as well.

Acknowledgment

This work was supported in part by TerraSwarm, one of six centers of STARnet, a Semiconductor Research Corporation program sponsored by MARCO and DARPA.

References

[1] W. J. Kop, R. J. Verdino, J. S. Gottdiener, S. T. O'Leary, C. Bairey Merz, and D. S. Krantz, "Changes in heart rate and heart rate variability before ambulatory ischemic events," *Journal of the American College of Cardiology*, vol. 38, no. 3, pp. 742–749, 2001.

[2] M. Raya and L. Sison, "Adaptive noise cancelling of motion artifact in stress ECG signals using accelerometer," in *Engineering in Medicine and Biology, 24th Annual Conference and the Annual Fall Meeting of the Biomedical Engineering Society EMBS/BMES Conference, 2002. Proceedings of the Second Joint*, vol. 2, pp. 1756–1757 vol.2, 2002.

[3] H. Kim, S. Kim, N. Van Helleputte, T. Berset, D. Geng, I. Romero, J. Penders, C. Van Hoof, and R. Yazicioglu, "Motion artifact removal using cascade adaptive filtering for ambulatory ECG monitoring system," in *Biomedical Circuits and Systems Conference (BioCAS), 2012 IEEE*, pp. 160–163, Nov 2012.

[4] K. Sweeney, T. Ward, and S. McLoone, "Artifact removal in physiological signals - practices and possibilities," in *Information Technology in Biomedicine, IEEE Transactions on*, vol. 16, pp. 488–500, May 2012.

[5] O. Sayadi and M. Shamsollahi, "ECG denoising and compression using a modified extended kalman filter structure," in *Biomedical Engineering, IEEE Transactions on*, vol. 55, pp. 2240–2248, Sept 2008.

[6] I. Romero, "PCA and ICA applied to noise reduction in multi-lead ECG," in *Computing in Cardiology, 2011*, pp. 613–616, Sept 2011.

[7] S. Kim, L. Holmstrom, and J. McNames, "Multiharmonic tracking using marginalized particle filters," in *Engineering in Medicine and Biology Society, 2008. EMBS 2008. 30th Annual International Conference of the IEEE*, pp. 29–33, Aug 2008.

[8] C. Lin, M. Bugallo, C. Mailhes, and J.-Y. Tourneret, "ECG denoising using a dynamical model and a marginalized particle filter," in *Signals, Systems and Computers (ASILOMAR), 2011 Conference Record of the Forty Fifth Asilomar Conference on*, pp. 1679–1683, Nov 2011.

[9] C. Lin, A. Giremus, C. Mailhes, and J.-Y. Tourneret, "Beat-to-beat p and t wave delineation in ECG signals using a marginalized particle filter," in *Signal Processing Conference (EUSIPCO), 2012 Proceedings of the 20th European*, pp. 479–483, Aug 2012.

[10] 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.

[11] 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.

[12] 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.

[13] 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.

[14] M. Arulampalam, S. Maskell, N. Gordon, and T. Clapp, "A tutorial on particle filters for online nonlinear/non-gaussian bayesian tracking," *Signal Processing, IEEE Transactions on*, vol. 50, pp. 174–188, Feb 2002.

[15] 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). Circulation Electronic Pages: http://circ.ahajournals.org/cgi/content/full/101/23/e215 PMID:1085218; doi: 10.1161/01.CIR.101.23.e215.

[16] 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.