Application of Filtering Methods for Removal of Resuscitation Artifacts from Human ECG Signals

Ivan Markovsky, Anton Amann, and Sabine Van Huffel

Abstract— Band-pass, Kalman, and adaptive filters are used for removal of resuscitation artifacts from human ECG signals. A database of separately recorded human ECG and animal resuscitation artifact signals is used for evaluation of the methods. The considered performance criterion is the signal-to-noise ratio (SNR) improvement, defined as the ratio of the SNRs of the filtered signal and the given ECG signal. The empirical results show that for low SNR of the given signal, a band-pass filter yields the best performance, while for high SNR, an adaptive filter yields the best performance.

I. INTRODUCTION

A measured ECG signal $y = (y(1), \ldots, y(T))$ has two components:

- the ventricular fibrillation ECG signal v (the useful signal) and
- the resuscitation artifacts c (the disturbance).

Our goal is to extract the unknown useful signal v from the given signal y. Often another signal u is given that has causal relation with the artifact c.

In general, different methods for artifacts removal produce different approximations $\widehat{v}$ of the unknown signal v. The quality of approximation of v by $\widehat{v}$ is evaluated in [6] by the signal-to-noise ratio (SNR)

$$\text{SNR}(\widehat{v}, v) := 20 \log_{10} \left(\frac{\|v\|}{\|v - \widehat{v}\|} \right), \quad (1)$$

where $\| \cdot \|$ denotes the 2-norm $\|v\| := \sqrt{\sum_{t=1}^{T} v^2(t)}$. Of course, $\text{SNR}(\widehat{v}, v)$ can not be computed in real-life applications. It is used for evaluation of the methods on artificially constructed ECG signals y, where the true signal v is known.

The empirical results published in [6] do not compare the proposed method with other methods from the literature. We fulfill this gap in the literature by presenting a comparative study of the methods of [6] with classical filtering methods. Section II described the filtering methods, used in the comparison. Section III presents the performance of the methods on a database of separately recorded resuscitation artifacts and reference signals from pigs and ventricular fibrillation ECG signals from human. The same data is used in [6] for tuning and evaluation of an adaptive filtering method.

I. Markovsky is with the School of Electronics and Computer Science, University of Southampton, SO17 1BJ, UK, im@ecs.soton.ac.uk

A. Amann is with Innsbruck Medical University and Department of Anesthesia and General Intensive Care, Anichstr 35, A-6020 Innsbruck, Austria, Anton.Amann@i-med.ac.at

S. Van Huffel is with K.U.Leuven, ESAT-SCD, Kasteelpark Arenberg 10, B-3001 Leuven, Belgium Sabine.VanHuffel@esat.kuleuven.be

II. METHODS FOR ARTIFACTS REMOVAL

A. Band-pass filtering

A key assumption for applying band-pass filtering is that the useful signal, the ventricular fibrillation ECG signal v, and the disturbance, the resuscitation artifacts c, have well separated spectral bands. The left plot in Figure 1 shows the power spectral densities of v and c in a particular experiment. The plot shows that up to 3Hz the spectrum of c dominates the spectrum of v and for frequencies above 3Hz the opposite is true—the spectrum of v dominates the spectrum of c. Therefore, low-pass filtering, with a cut-off frequency $f_0 = 3Hz$, can extract from the signal $y = c + v$ an approximation $\widehat{c}$ of the resuscitation artifacts c and then find an approximation of v as $\widehat{v} := y - \widehat{c}$.

The structure of the low-pass filter for artifacts removal is shown in Figure 2. Note that this method does not use the reference signal u, i.e., the method can be applied even in the case when no reference signal is available.


Fig. 2. Structure of the low-pass filter for ECG artifacts removal. The filter F is low-pass with a cut-off frequency f_0

We design a finite impulse response (FIR) low-pass filter with 100 time lags using the window method [5]. The ideal low-pass filter has magnitude one at all frequencies with magnitude less than the cut-off frequency and magnitude zero at all other frequencies. The impulse response of this ideal filter is infinite and non-causal, so that it is not implementable. The windowing method resolves the implementability problem by truncating the ideal impulse response by a window of a specified length. In the experiments, we use the Hamming window [5].

B. Kalman filtering

By assumption the reference signal u has a causal relation with the artifact signal c. Formally, this means that there is a model M, such that when u is given as an input to M and the initial conditions are properly chosen, the resulting output is c. The model M, however, may be a complicated nonlinear time-varying one, so that we consider simple linear time-invariant models. In addition, apart from u other unmeasurable signals may affect c. Finally, there may be measurement noise on the reference and the corrupted ECG signals. For these reasons, in practise, we can not obtain

Fig. 1. Power spectral densities (PSD) of v and c (left) and magnitude response of the low-pass filter (right).

an exact relation between u and c. The model M is only an approximation of the unknown exact relation $u \mapsto c$.

An artifact removal method, called "naive model based method", using a model M for the relation between u and c is shown on the left plot of Figure 3. This method uses the reference signal u and the prior knowledge about the relation between u and c in the form of the model M. Note that the naive method sets the initial conditions for M to zero. This is an arbitrary choice (hence the name of the method).

A proper way of choosing the initial conditions is actually given by the celebrated Kalman filter. The Kalman filter for M accepts as inputs both u and y and recursively updates the initial conditions for optimal prediction $\widehat{c}$ of c. The Kalman filter, however, uses extra prior knowledge: an initial value x_{ini} for the initial conditions and a matrix P_{ini} related to the uncertainty of (or, equivalently, the confidence in) x_{ini}. In addition, the Kalman filter produces optimal prediction of c in certain specified sense, that involves choice of an approximation criterion.


Fig. 3. Left: naive model based method for artifact removal. The initial conditions of the model M are set to zero. Right: method for artifact removal based on the Kalman filter F for M. The initial conditions are recursively updated using the data (u, y), starting from an initial guess x_{ini} with uncertainty P_{ini}.

The main question in applying the naive model based method or the optimal (Kalman) filtering method is the selection of the model M. In a practical artifacts removal problem such a model is *not* given as a prior knowledge but has to be deduced from the data. This lead us to the adaptive filtering methods that compute on-line a model M and, based on M, filter the signal. Before explaining the adaptive methods, however, we give some background information on offline identification of a model from data.

We adopt the deterministic identification setting of [4],

where the aim is to minimise the fitting error $\|y - \widehat{c}\|$ over the model parameters and the initial conditions. The criterion $\|y - \widehat{c}\|$ corresponds to what is called output error identification [3], however, no stability constraint is imposed on the identified model.

We consider a linear time-invariant model with a state space representation

$$\begin{aligned} x(t+1) &= Ax(t) + Bu(t) \\ \widehat{c}(t) &= Cx(t) + Du(t). \end{aligned} \tag{2}$$

The Kalman filter is a linear system derived from the identified model, the approximation criterion, and the assumption about the initial conditions $(x_{\text{ini}}, P_{\text{ini}})$. The computational load of finding the approximation $\widehat{c}$ by the Kalman filter depends on the order of the model.

C. Adaptive filtering

Adaptive filters are conceptually similar to the Kalman filter: they are model-based and perform optimal least squares filtering. The main difference is that in the classical Kalman filter the model is identified off-line or is given as prior knowledge, while in adaptive filters the model is identified from the data in real-time. The recursive identification procedure is the essential part of any adaptive filter.

There is a large variety of adaptive filters depending on the model class and the identification algorithm that are used. A common model class is the set of FIR models with at most l lags

$$\widehat{c}(t) = \widehat{h}_0(t)u(t) + \widehat{h}_1(t)u(t-1) + \cdots + \widehat{h}_l(t)u(t-l). \tag{3}$$

For a time-invariant FIR model (i.e., $\widehat{h}$ constant in time), l specifies the model complexity. The vector

$$\widehat{h}(t) := \text{col}\left(\widehat{h}_0(t), \widehat{h}_1(t), \ldots, \widehat{h}_l(t)\right)$$

is the parameter of the FIR model at time t. At each time step, $\widehat{h}(t)$ is determined as a minimum point of a certain cost function depending on $\widehat{h}(t)$ and the data (u, y) (up to and including time instant t, due to the causality requirement). A

typical cost function in adaptive filtering is the 2-norm of the error signal $e := y - \widehat{c}$ over the window $t - t_1, t - t_1 + 1, \ldots, t$

$$J\big(\widehat{h}(t)\big) := \sum_{\tau=0}^{t_1} \big(y(t-\tau) - \widehat{c}(t-\tau)\big)^2. \tag{4}$$

Common algorithms for adaptive filtering that minimise J are the recursive least squares and gradient descent algorithms. A class of adaptive filters, using gradient descent, is the one of least mean squares (LMS) adaptive filters [1]. The parameter $\widehat{h}$ is selected adaptively from the data, while the parameter l is fixed during the operation of the algorithm and is user defined. In order to distinguish between these two types of parameters, we call the latter ones *hyper-parameters*. The window length t_1 is a hyper-parameter that determines the speed of adaptation (larger values of t_1 imply slower adaptation).

The cost function J, forces the approximation $\widehat{c}$ to be as close as possible to the measurement y in a finite horizon 2-norm sense (specified by the hyper-parameter t_1). The adaptive filters determine the model parameters $\widehat{h}(t)$ as a minimum point of J and obtain the filtered signal $\widehat{c}(t)$ at time t from the identified model (3). An important implicit assumption, on which the adaptive methods are based, is that by minimizing J, the filtered signal $\widehat{c}$ becomes closer to the artifact signal c.

The method of [6] is based on additional assumptions that there is a "true" FIR model

$$
\begin{aligned}
c(t) &= \bar{h}_0(t)u(t) + \bar{h}_1(t)u(t-1) + \cdots + \bar{h}_l(t)u(t-l) \\
&=: \underbrace{\big[u(t) \quad u(t-1) \quad \cdots \quad u(t-l)\big]}_{\mathbf{u}(t)} \bar{h}(t),
\end{aligned} \tag{5}
$$

i.e., the artifact signal c satisfies an FIR model with l taps for certain "true" parameter $\bar{h}$ and that $\bar{h}$ satisfies the "random walk" equation

$$\bar{h}(t+1) = \bar{h}(t) + w(t), \tag{6}$$

where w is a zero-mean stationary white stochastic process. Finally, it is assumed that the ECG signal $v = y - c$ is a zero-mean stationary white stochastic process and is independent of w. Under these assumptions, (5) and (6) form a state-space representation of a classical linear time-varying stochastic system

$$
\begin{aligned}
\bar{h}(t+1) &= \bar{h}(t) + w(t) \\
y(t) &= \underbrace{\mathbf{u}(t)\bar{h}(t)}_{c(t)} + v(t).
\end{aligned} \tag{7}
$$

The parameters of this model are *known*, so the adaptive filtering problem reduces to the simpler and easier linear state estimation problem, which optimal (in the minimum variance sense) solution is the Kalman filter. Note that the identification of (5) is implicitly done by the Kalman filter for (7).

Let W be the covariance matrix of $w(0)$ and let V be the covariance matrix of $v(0)$. Under the above stated assumptions for v and w, the Kalman filter for (7) is

$$\widehat{h}(t+1) = \widehat{h}(t) + \Sigma(t)\mathbf{u}^\top(t)\big(\mathbf{u}(t)\Sigma(t)\mathbf{u}^\top(t) + V\big)^{-1}\big(y(t) - \mathbf{u}\widehat{h}(t)\big)$$
$$\Sigma(t+1) = \Sigma(t) + W - \Sigma(t)\mathbf{u}^\top(t)\big(\mathbf{u}(t)\Sigma(t)\mathbf{u}^\top(t) + V\big)^{-1}\mathbf{u}(t)\Sigma(t).$$

with $\widehat{h}(0) = h_{\text{ini}}$ and $\Sigma(0) = \Sigma_{\text{ini}}$. It defines a minimum variance estimator $\widehat{h}$ for $\bar{h}$ and the corresponding $\widehat{c} = \mathbf{u}(t)\widehat{h}(t)$ is the minimum variance estimate of c. Hyper-parameters in this case are the filter length l, the covariance matrices W and V, and the initial conditions h_{ini}, Σ_{ini}, and $u(0), u(-1), \ldots, u(1-l)$.

The stochastic assumptions imposed on w and v are motivated by the reformulation of the original adaptive filtering problem as a classical linear state estimation problem, of which recursive solution is known. These assumptions, however, are hard to interpret in the context of the original ECG artifacts removal problem and their relevance for the real-life problem is unclear.

III. RESULTS: PERFORMANCE EVALUATION

The data used in the performance evaluation is obtained at the Innsbruck Medical University and consists of seven ventricular fibrillation ECG signals v and reference signals u (the arterial blood pressure) from human and the seven resuscitation artifact signals c from pigs. We combined the ventricular fibrillation and resuscitation signals pairwise to form 49 test signals y. The evaluated methods are applied on the 49 signals and the 49 SNRs of the estimated signals $\widehat{v}$ are computed. Table I and Figure 4 show the average results for SNR of the signal y, -10, -5, 0, 5, and 10. The standard deviations (over the 49 experiments) are is given in brackets in the table and is visualised with the vertical bars on the figure.

For $\text{SNR}(y, v)$ less than -3 dB, the low-pass filter achieves the best performance. A possible explanation is that for low SNR, the system identification methods fail to obtain a sufficiently good model for the observed data and the Kalman filter, which is based on the model, is sensitive to model-data discrepancy. Indeed, obtaining a good model (2) for the artifact signal c is a challenging problem even when the data used for identification is (u, c) (reference method). Note that the signals that we use in the simulation study are measured in a real-life environment and therefore they need not satisfy a linear time-invariant model (2) of low order. Moreover, the actual identification problem encountered in ECG artifact removal is to derive the model from (u, y) in real-time.

As expected, the reference method outperforms the Kalman filtering method. The explanation is that the reference method is noncausal and in the system identification step uses information (the true artifact signal c) that is not available to the other methods. The Kalman filtering method is also noncausal and therefore has an advantage over the adaptive filtering method, however, it is based on a time-invariant model, which is a limitation when the data is generated by a time-varying "true" system. For high SNR the best performance is achieved by the adaptive filtering method of [6], which suggests that the true model for v is indeed time-varying.

TABLE I

AVERAGE PERFORMANCE $\mathrm{SNR}(\widehat{v}, v)$ IN dB OF THE FILTERING METHOD. IN BRACKETS IS THE STANDARD DEVIATION OVER THE 49 EXPERIMENTS.

Method	−10	−5	0	5	10
Low-pass	1.5 (2.4)	4.8 (2.0)	6.8 (1.8)	7.8 (1.9)	8.2 (1.9)
Kalman	−2.5 (1.9)	2.0 (1.8)	6.2 (1.8)	9.2 (1.9)	11.2 (2.9)
Reference	−1.3 (1.8)	3.7 (1.8)	8.7 (1.8)	13.7 (1.8)	18.7 (1.8)
LMS	−1.5 (2.0)	2.7 (1.8)	6.0 (1.6)	8.0 (1.6)	9.0 (1.6)
Rheinberger et al.	−5.2 (1.4)	−0.2 (1.4)	4.6 (1.4)	9.2 (1.3)	13.0 (1.3)
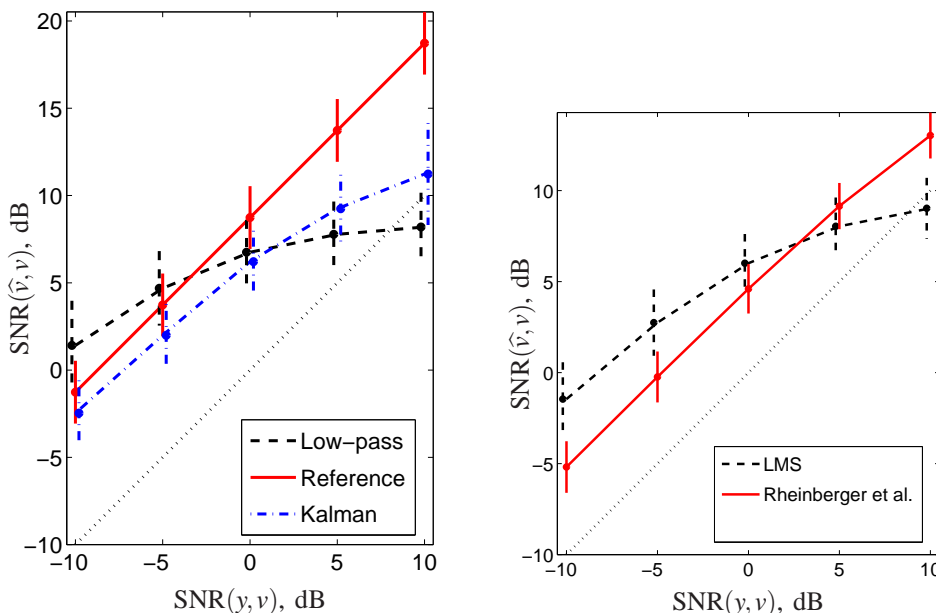
Fig. 4. Comparison of the methods in terms of average SNR improvement. Left figure: nonadaptive methods, right figure: adaptive methods. The "reference" method is the Kalman filter based on a model identified from the data (u, c) while the method labeled with "Kalman" is the Kalman filter based on a model identified from the data (u, y).

IV. CONCLUSIONS

The robustness of the filtering methods, i.e., their sensitivity to the models being used, is crucial in the removal of resuscitation artifacts from ventricular fibrillation ECG signals. The low-pass filter uses the simplest model: separation of the spectra of the useful signal and the disturbance. The only tuning parameter in this case is the cut-off frequency. Because of its simplicity the low-pass filtering method is more robust than the Kalman filter, which is based on an accurate model of the data. The experiments show, however, that even in the unrealistic case when the model, used in the synthesis of the Kalman filter, is identified from the unknown true data, the low-pass filter on the average still achieves better performance than the Kalman filter, provided the SNR of the given ECG signal is sufficiently low. For high SNR the best performance achieves the adaptive method of [6].

ACKNOWLEDGMENTS

REFERENCES

[1] S. Haykin. *Adaptive filter theory*. Prentice Hall, 1991.
[2] J. Husøy, J. Eilevstrønn, T. Eftestøl, S. Aase, H. Myklebust, and P. Steen. Removal of cardiopulmonary resuscitation artifacts from human ECG using an efficient matching pursuit-like algorithm. *IEEE Trans. Biomed. Eng.*, 49:1287–1298, 2002.
[3] L. Ljung. *System Identification: Theory for the User*. Prentice-Hall, Upper Saddle River, NJ, 1999.
[4] I. Markovsky, J. C. Willems, S. Van Huffel, and B. De Moor. *Exact and Approximate Modeling of Linear Systems: A Behavioral Approach*. SIAM, March 2006.
[5] A. Oppenheim and A. Willsky. *Signals and Systems*. Prentice Hall, 1996.
[6] K. Rheinberger, T. Steinberger, K. Unterkofler, M. Baubin, A. Klotz, and A. Amann. Removal of CPR artifacts from the ventricular fibrillation ECG by adaptive regression on lagged reference signals. *IEEE Trans. Biomed. Eng.*, 55(1):130–137, 2008.