Joint Denoising and Narrowband Artifact Rejection for ECG Signals

Antonio Fasano, Valeria Villani

Università Campus Bio-Medico di Roma, Rome, Italy

Abstract

ECG signals are corrupted by various kinds of noise and artifacts that may negatively affect any subsequent analysis. In particular, narrowband artifacts include power-line interference and harmonic artifacts. Customarily, noise reduction and artifact rejection are tackled as two distinct problems. In this paper, we propose a joint *approach to denoising and narrowband artifact rejection that exploits the local structure of a noisy ECG. Simulation results confirm the effectiveness of the approach and highlight a notable ability to remove both noise and narrowband artifacts in ECG signals.*

1. Introduction

The electrocardiogram (ECG) is a non invasive measure of the electrical activity of the heart recorded by skin electrodes. Unfortunately, ECG signals are highly susceptible to electromagnetic interferences and can be contaminated by several kinds of noise and artifacts, such as 50 or 60 Hz power-line noise, electromyographic noise, baseline wander, and measurement noise [1]. In particular, narrowband artifacts include harmonic artifacts and power-line interference, which is ubiquitous in clinical environment [1]. It is caused by supply plugs and cables, and sometimes can mask ECG signals, especially those portions having low amplitude [1, 2]. Although modern biomedical amplifiers have a very high common mode rejection ratio, recordings are often contaminated by residual power-line interference [3, 4]. Thus, denoising and artifact suppression are unavoidable steps in any processing of ECG signals before clinical interpretation. However, particular care must be paid to preserve signal morphology and retain all relevant clinical information.

Customarily, noise reduction and artifact rejection are tackled as two distinct problems. Classical approaches resort to linear time-invariant filters, whether low-pass to attenuate noise, or notch to remove power-line interference [1, 2, 4]. However, low-pass filtering cannot remove in-band noise due to muscle activity, nor accommodate the non-stationary nature of cardiac signals [2]. Moreover, since power-line amplitude and phase are time-varying, linear time-invariant notch filters are unable to track slow variations around the $50/60$ Hz fundamental and its harmonics.

In the recent work [5], we proposed a novel approach to smoothing and denoising ECG records. Our algorithm is based on the notion of quadratic variation and keeps into account the local structure of the ECG signal. It proved being quite good in smoothing ECGs heavily affected by noise, but its effectiveness reduces when it is used to reject multiple harmonic artifacts, like power-line interference. To overcome this limitation, in this paper we propose a *joint* approach to denoising and narrowband artifact rejection that exploits the local structure of a noisy ECG.

The paper is organized as follows. The rationale behind the proposed approach is described in Section 2. The smoothing algorithm is derived in Section 3. Sections 4 and 5 follow with simulation results and conclusions.

2. Rationale

2.1. The quadratic variation

Measured ECG is affected by several kinds of noise and artifacts whose effect is to introduce additional "variability" into the observed ECG with respect to the true one. The amount of additional "variability" is time-varying, since signal morphology changes over time, and noise and artifacts can be persistent, like thermal noise, or transient, like electromyographic noise. As a consequence, different portions of a noisy ECG present different local SNRs, where local SNR is the SNR computed on a portion of the signal. Thus, provided that we introduce a suitable index of variability, smoothing and denoising can be performed by locally reducing the variability of the measured ECG. This means reducing the variability of different portions of a noisy ECG by an amount that is inversely related to the local SNR. The variability of a generic vector can be quantified introducing the following

Definition 1: Given a vector $\boldsymbol{x} = [x_1 \cdots x_n]^{\mathrm{T}} \in \mathbb{R}^n$, the *quadratic variation* of $\boldsymbol{x}$, namely $[\boldsymbol{x}]$, is defined as

$$[\boldsymbol{x}] \doteq \sum_{k=1}^{n-1} (x_k - x_{k+1})^2 = \|\boldsymbol{D}\boldsymbol{x}\|^2 \qquad (1)$$

where $\|\cdot\|$ denotes the Euclidean norm and $\boldsymbol{D}$ is the $(n-1)\times n$ matrix

$$\boldsymbol{D} = \begin{bmatrix} 1 & -1 & & \\ & \ddots & \ddots & \\ & & 1 & -1 \end{bmatrix}. \qquad (2)$$

Let us denote by $\boldsymbol{x}(k_i, k_{i+1}) = \begin{bmatrix} x_{k_i} \cdots x_{k_{i+1}} \end{bmatrix}^{\mathrm{T}}$, with $1 \le k_i < k_{i+1} \le n$, a generic subvector of $\boldsymbol{x}$. Introducing the diagonal matrix $\boldsymbol{S}(k_i, k_{i+1}) \in \mathbb{R}^{n\times n}$ having the diagonal entries with indices $k_i, \ldots, k_{i+1}$ equal to one and the other entries zero, the quadratic variation of the segment $\boldsymbol{x}(k_i, k_{i+1})$ can be expressed in terms of $\boldsymbol{x}$ as

$$[\boldsymbol{x}(k_i, k_{i+1})] = \|\boldsymbol{D}\boldsymbol{S}(k_i, k_{i+1})\boldsymbol{x}\|^2 = \|\boldsymbol{\mathfrak{D}}_i\boldsymbol{x}\|^2 \qquad (3)$$

where $\boldsymbol{\mathfrak{D}}_i = \boldsymbol{D}\boldsymbol{S}(k_i, k_{i+1})$.

In [5] we proved that the quadratic variation is a consistent index of variability for vectors, and we motivated its use in denoising signals. However, considering a typical scenario, let $\boldsymbol{x} = \boldsymbol{x}_0 + \boldsymbol{m} + \boldsymbol{a}$, where $\boldsymbol{x}_0$ is a deterministic vector, $\boldsymbol{m}$ and $\boldsymbol{a}$ are independent zero-mean random vectors due to measurement noise and harmonic artifacts, e.g., $50/60$ Hz power-line noise. We may assume $\boldsymbol{m} \sim \mathcal{N}\left(\boldsymbol{0}, \sigma_m^2\boldsymbol{I}\right)$ and $\boldsymbol{a}$ vector of samples from a harmonic process with random amplitude A, and frequency f_0. It is easy to verify [5] that the average quadratic variation of $\boldsymbol{x}$ can be expressed as

$$\mathbb{E}\{[\boldsymbol{x}]\} = [\boldsymbol{x}_0] + K\left(\mathrm{SNR}^{-1} + \alpha\mathrm{SIR}^{-1}\right) \qquad (4)$$

where K is a positive constant, $\alpha = 4\sin^2(\pi f_0/F_c)$, with F_c the sampling frequency, $\mathrm{SNR} = {}^{\|\boldsymbol{x}_0\|^2}\!/_{n\sigma_m^2}$ denotes the signal-to-noise ratio and $\mathrm{SIR} = {}^{2\|\boldsymbol{x}_0\|^2}\!/_{n\mathbb{E}\{A^2\}}$ is the signal-to-interference ratio, considering the power-line noise as interference. From (4) it is evident that the average quadratic variation is a decreasing function of SNR and SIR. This supports the rationale behind quadratic variation reduction as a viable criterion for noise reduction [5]. However, the coefficient α multiplying SIR^{-1} is less than 1 when $f_0 < F_c/6$, and decreases as the ratio f_0/F_c decreases. For example, when $f_0 = 50$ Hz and $F_c = 1024$ Hz, the coefficient $\alpha \simeq 9\times 10^{-2}$. As a result, low-frequency harmonic artifacts tend to be less attenuated in response to a quadratic variation reduction. Thus, extra conditions must be considered for narrowband artifact rejection.

2.2. Narrowband artifact rejection

The approach proposed to effectively reject narrowband artifacts is to exploit quadratic variation reduction *in conjunction* with an additional requirement: to make negligible the energy content of narrowband artifacts in the denoised signal. To quantify it, let $\boldsymbol{X} = \boldsymbol{W}\boldsymbol{x}$ be the DFT [6] of $\boldsymbol{x}$, where $\boldsymbol{W}$ is the DFT matrix. Now, denote by $\widetilde{\boldsymbol{W}}$ the matrix obtained stacking the rows of $\boldsymbol{W}$ corresponding to the harmonic components of the narrowband artifacts that we want to reject. Note that since $\boldsymbol{x}$ is a real vector, symmetries occur in its DFT [6] and rows have to be matched in pairs in general. Matrix $\widetilde{\boldsymbol{W}}$ has dimensions $m\times n$, with $m < n$ in general. The quadratic form

$$\mathcal{E}_a = \left\|\widetilde{\boldsymbol{W}}\boldsymbol{x}\right\|^2 = \boldsymbol{x}^{\mathrm{T}}\mathrm{Re}\left\{\widetilde{\boldsymbol{W}}^{\mathrm{H}}\widetilde{\boldsymbol{W}}\right\}\boldsymbol{x} \qquad (5)$$

quantifies the energy content of narrowband artifacts, where $(\cdot)^{\mathrm{H}}$ denotes the transpose conjugate and $\mathrm{Re}\{\cdot\}$ the real part.

In the next section we exploit these results and develop an effective algorithm for smoothing ECG signals.

3. Smoothing ECG signals

In this section, we denote by $\boldsymbol{q}$ the vector collecting n samples of the measured ECG, which is affected by noise and artifacts, and by $\boldsymbol{x}$ the corresponding vector after smoothing. The idea is to perform smoothing, meant as combined denoising and artifact rejection, by *jointly* reducing the local quadratic variation of different segments of the measured ECG and the energy content of narrowband artifacts.

Let the vector $\boldsymbol{q}$ be decomposed into $L+1$ segments

$$\boldsymbol{q}(k_i, k_{i+1}) = [q_{k_i} \cdots q_{k_{i+1}}]^{\mathrm{T}}, \quad \text{for } i = 0, \ldots, L \qquad (6)$$

with $0 \le L \le n-1$ and $1 = k_0 < k_1 < \cdots < k_L < k_{L+1} = n$. Segments in (6) denote distinct portions of the ECG characterized by different local SNRs, like QRS complexes, P-waves or T-waves. Note that two consecutive segments overlap. This choice guarantees the absence of abrupt changes in the smoothed vector $\boldsymbol{x}$.

Following the line of reasoning outlined above, denoising and narrowband artifact rejection can be achieved by solving the following convex optimization problem

$$\begin{cases} \underset{\boldsymbol{x}\in\mathbb{R}^n}{\text{minimize}} & \|\boldsymbol{x}-\boldsymbol{q}\|^2 \\ \text{subject to} & [\boldsymbol{x}(k_i, k_{i+1})] \le a_i,\ i = 0, \ldots, L \\ & \mathcal{E}_a \le b \end{cases} \qquad (7)$$

where $a_0, \ldots, a_L$ are positive constants controlling the degree of smoothness applied to each segment $\boldsymbol{q}(k_i, k_{i+1})$ of $\boldsymbol{q}$, whereas b is a positive constant controlling the degree of rejection of narrowband artifacts. We do not need to know in advance the appropriate values for a_i and b in any particular problem.

It is possible to prove that the solution to (7) is given by

$$\boldsymbol{x} = \left(\boldsymbol{I} + \sum_{i=0}^{L}\lambda_i\boldsymbol{\mathfrak{D}}_i^{\mathrm{T}}\boldsymbol{\mathfrak{D}}_i + \nu\mathrm{Re}\left\{\widetilde{\boldsymbol{W}}^{\mathrm{H}}\widetilde{\boldsymbol{W}}\right\}\right)^{-1}\boldsymbol{q} \qquad (8)$$

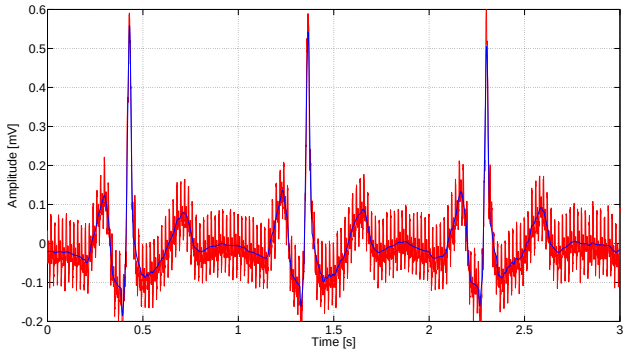

Figure 1. ECG from real data: noisy ECG $\boldsymbol{q}$ (red) and smoothed ECG $\boldsymbol{x}$ (blue) using the proposed algorithm.

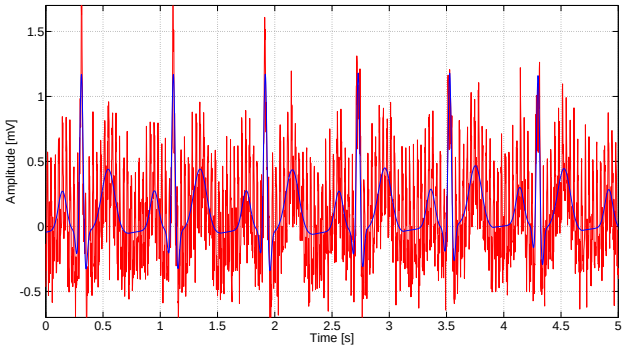

Figure 2. Synthetic ECG: noiseless reference ECG $\boldsymbol{q}_0$ (blue) and corrupted ECG $\boldsymbol{q}$ (red).

where $\boldsymbol{I}$ denotes the identity matrix, and $\lambda_0, \ldots, \lambda_L$ and ν are nonnegative parameters that control the degree of smoothing applied to $\boldsymbol{q}$. These parameters are related to a_i and b in (7), but are used in their place to control the solution (8). Smoothing can be performed without caring about a_i and b, by adapting λ_i and ν in (8) to meet some performance criteria. It is worthwhile to remark that, even though $\lambda_0, \ldots, \lambda_L$ and ν interact, λ_i mainly controls the quadratic variation of the segment $\boldsymbol{x}(k_i, k_{i+1})$, whereas ν mainly controls the degree of rejection of narrowband artifacts.

Finally, some remarks on computational aspects, since matrix inversion is involved in (8). Let us denote by n the size of $\boldsymbol{q}$. In the general case $\lambda_i > 0$ and $\nu > 0$, it is possible to prove that smoothing through (8) can be performed with complexity $O(n \log n)$ using the conjugate gradient method [7]. When $\lambda_i = 0$, for $i = 0, \ldots, L$ and $\nu > 0$, the complexity is still $O(n \log n)$, but the solution is achieved by direct methods. Eventually, when $\nu = 0$, (8) can be solved efficiently by direct methods with complexity $O(n)$ *regardless* of L.

4. Simulation results

The performance of the proposed algorithm has been investigated both on real and synthetic ECG signals. In all simulations, ECG delineation was performed manually: the different segments were roughly determined by visual inspection. No accurate delineation is required, as the proposed algorithm proved to be robust in this regard.

Regarding real signals, in Figure 1 we report a portion of a noisy ECG $\boldsymbol{q}$ (red) strongly affected by narrowband artifacts centered at $50\,\text{Hz}$ and $100\,\text{Hz}$, due to power-line noise. In the same figure we report the corresponding smoothed signal $\boldsymbol{x}$ (blue) resulting from joint denoising and narrowband artifact rejection. Although the original ECG was highly corrupted, the proposed approach managed to effectively smooth it.

In order to provide quantitative evaluation of performances, we tested the proposed algorithm on synthetic ECG signals corrupted by noise and harmonic artifacts. A synthetic noiseless ECG record, denoted by $\boldsymbol{q}_0$, was generated according to the model described in [8], setting the sampling frequency to $512\,\text{Hz}$ and the heart rate to $75\,\text{bpm}$. The bandwidth of $\boldsymbol{q}_0$ essentially does not exceed $40\,\text{Hz}$. The noiseless ECG $\boldsymbol{q}_0$ was corrupted by additive noise, denoted by $\boldsymbol{w}$, which is a zero-mean Gaussian random vector with independent components having variance σ_w^2 such that $\text{SNR} = 10 \log \frac{\|\boldsymbol{q}_0\|^2}{n \cdot \sigma_w^2} = 5\,\text{dB}$, being n the length of $\boldsymbol{q}_0$. In addition, $\boldsymbol{q}_0$ was corrupted by harmonic artifacts, denoted by $\boldsymbol{d}$, which consist of three sine waves with random phases and frequencies $30\,\text{Hz}$, $60\,\text{Hz}$ and $120\,\text{Hz}$, respectively. The corresponding amplitudes, namely A_{30}, A_{60} and A_{120}, are in the ratio $A_{30}/A_{60} = A_{60}/A_{120} = 2$ and are such that the corresponding signal-to-interference ratio $\text{SIR} = 10 \log \frac{\|\boldsymbol{q}_0\|^2}{\|\boldsymbol{d}\|^2} = 0\,\text{dB}$. The sine waves at $60\,\text{Hz}$ and $120\,\text{Hz}$ account for the first and second harmonics of the power-line noise, whereas the sine wave at $30\,\text{Hz}$ models a strong in-band harmonic interference.

The resulting ECG corrupted by noise and artifacts is $\boldsymbol{q} = \boldsymbol{q}_0 + \boldsymbol{w} + \boldsymbol{d}$, and is characterized by a signal-to-noise-plus-interference ratio $\text{SNIR}_0 = 10 \log \frac{\|\boldsymbol{q}_0\|^2}{\|\boldsymbol{w}+\boldsymbol{d}\|^2} = -1.3\,\text{dB}$. Note that $\boldsymbol{q}$ is highly corrupted. Figure 2 reports a segment of $\boldsymbol{q}$ (red) together with the corresponding noiseless reference ECG $\boldsymbol{q}_0$ (blue).

Performance is measured in terms of SNIR gain

$$\text{G}_{\text{SNIR}} = \text{SNIR}_\text{s} - \text{SNIR}_0 = 10 \log \frac{\|\boldsymbol{w} + \boldsymbol{d}\|^2}{\|\boldsymbol{x} - \boldsymbol{q}_0\|^2} \qquad (9)$$

where $\text{SNIR}_\text{s} = 10 \log \frac{\|\boldsymbol{q}_0\|^2}{\|\boldsymbol{x}-\boldsymbol{q}_0\|^2}$ is the signal-to-noise-plus-interference ratio after smoothing. By $(\boldsymbol{x} - \boldsymbol{q}_0)$ we regard as noise affecting $\boldsymbol{x}$ both the residual noise and the reconstruction error.

In Figure 3 we report the reference ECG $\boldsymbol{q}_0$ (blue)

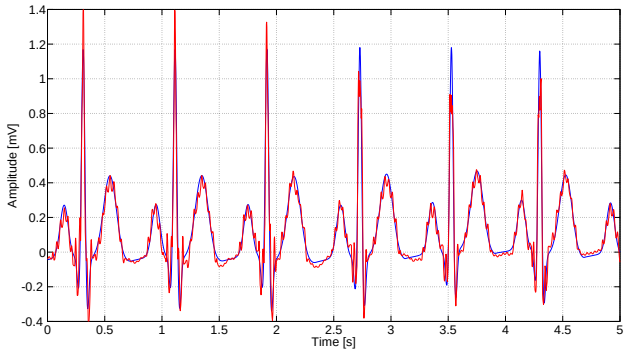

Figure 3. Reference ECG $\boldsymbol{q}_0$ (blue) and reconstructed ECG $\widetilde{\boldsymbol{x}}$ (red) after partial smoothing, i.e., $\nu = 0$ in (8).

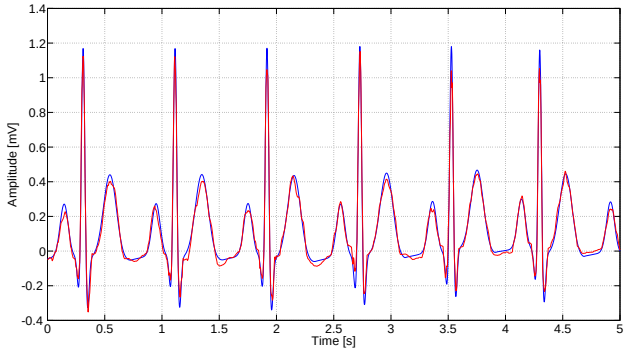

Figure 4. Reference ECG $\boldsymbol{q}_0$ (blue) and smoothed ECG $\boldsymbol{x}$ (red) after joint denoising and artifact rejection.

and the reconstructed record (red) resulting from partial smoothing, denoted by $\widetilde{\boldsymbol{x}}$, which reduces the quadratic variation of the noisy signal but not the energy content of harmonic artifacts, i.e., $\nu = 0$ in (8). Smoothing was applied with different smoothing parameters to the following segments: isoelectric PQ, ST and TP segments (λ_{iso}), P-waves (λ_{P}), Q-waves (λ_{Q}), R peaks (λ_{R}), S-waves (λ_{S}), and T-waves (λ_{T}). The corresponding smoothing parameters were: $\lambda_{\text{iso}} = 8\lambda_{\text{P}}$, $\lambda_{\text{Q}} = \lambda_{\text{S}} = 0.2\lambda_{\text{P}}$, $\lambda_{\text{R}} = 3$, and $\lambda_{\text{T}} = \lambda_{\text{P}} = 60$. Note that smoothing parameters are related and tend to have approximately a constant ratio, as highlighted in [5]. This allows us to reduce the number of smoothing parameters. It is remarkable that even without optimization the smoothing performance is significant, as Figure 3 shows. Indeed, for these realizations of noise and narrowband artifacts the gain is $\text{G}_{\text{SNIR}} = 13.5\,\text{dB}$.

However, a residual oscillatory behavior is still present in $\widetilde{\boldsymbol{x}}$. To suppress it, we must consider the combined action of $\lambda_i > 0$, $i \in \{\text{iso}, \text{P}, \text{Q}, \text{R}, \text{S}, \text{T}\}$, and $\nu > 0$ in (8). This amounts to perform *joint* denoising and harmonic artifact rejection on $\boldsymbol{q}$. In Figure 4 we compare the resulting smoothed ECG signal $\boldsymbol{x}$ (red) and the noiseless ECG record $\boldsymbol{q}_0$ (blue) as reference. The smoothing parameters λ_i are the same as in Figure 3 and $\nu = 10^3$ (coarsely set). As Figure 4 highlights, the proposed approach managed to effectively smooth the ECG signal. The corresponding SNIR gain is as high as $\text{G}_{\text{SNIR}} = 19.6\,\text{dB}$, thus entailing an additional $6.1\,\text{dB}$ gain over the sole denoising of Figure 3.

5. Conclusions

In this work we considered the problem of noise and narrowband artifact removal in ECG signals. The proposed approach achieves denoising and narrowband artifact rejection by combining two actions: the local reduction of the quadratic variation of different segments of the ECG by an amount inversely related to the local SNR, and the global reduction of the energy content of narrowband artifacts. These actions are performed *jointly*. Simulation results confirm the effectiveness of the approach and highlight its ability to remove noise and artifacts in ECG signals. Finally, it is worthwhile noting that the proposed algorithm is not limited to ECG and can be effectively applied to a broader class of signals.

References

[1] Clifford GD, Azuaje F, McSharry P (eds.). Advanced Methods and Tools for ECG Data Analysis. Artech House, Inc., 2006.

[2] Sörnmo L, Laguna P. Bioelectrical Signal Processing in Cardiac and Neurological Applications. Academic Press, 2005.

[3] Levkov C, Mihov G, Ivanov R, Daskalov I, Christov I, Dotsinsky I. Removal of power-line interference from the ECG: a review of the subtraction procedure. BioMed Eng OnLine 2005;4(50).

[4] Gregg RE, Zhou SH, Lindauer JM, Helfenbein ED, Giuliano KK. What is inside the electrocardiograph? J Electrocardiol 2008;41:8–14.

[5] Fasano A, Villani V, Vollero L. ECG smoothing and denoising by local quadratic variation reduction. In International Symposium Applied Sciences in Biomedical and Communication Technologies, ISABEL '11. ACM, 2011; 64:1–5.

[6] Oppenheim AV, Schafer RW, Buck JR. Discrete-Time Signal Processing. 2nd edition. Prentice-Hall, Inc., 1999.

[7] Kelley CT. Iterative Methods for Linear and Nonlinear Equations, volume 16 of Frontiers in Applied Mathematics. SIAM, 1995.

[8] McSharry PE, Clifford GD, Tarassenko L, Smith LA. A dynamical model for generating synthetic electrocardiogram signals. IEEE Trans Biomed Eng 2003;50(3):289–294.

Address for correspondence:

Antonio Fasano
Università Campus Bio-Medico di Roma
via Álvaro del Portillo, 21
00128 Rome, Italy
E-mail: a.fasano@unicampus.it