

ECG P-WAVE SMOOTHING AND DENOISING BY QUADRATIC VARIATION REDUCTION

Antonio Fasano¹, Valeria Villani^{1,2}, Luca Vollero¹ and Federica Censi²

¹*Università Campus Bio-Medico di Roma, Rome, Italy*

²*Department of Technology and Health, Italian National Institute of Health, Rome, Italy*

Keywords: ECG, P-wave, Atrial fibrillation, Smoothing, Denoising, Quadratic variation, Convex optimization.

Abstract: Atrial fibrillation is the most common persistent cardiac arrhythmia and it is characterized by a disorganized atrial electrical activity. Its occurrence can be detected, and even predicted, through P-waves time-domain and morphological analysis in ECG tracings. Given the low signal-to-noise ratio associated to P-waves, such analysis are possible if noise and artifacts are effectively filtered out from P-waves. In this paper a novel smoothing and denoising algorithm for P-waves is proposed. The algorithm is solution to a convex optimization problem. Smoothing and denoising are achieved reducing the quadratic variation of the measured P-waves. Simulation results confirm the effectiveness of the approach and show that the proposed algorithm is remarkably good at smoothing and denoising P-waves. The achieved SNR gain exceeds 15 dB for input SNR below 6 dB. Moreover the proposed algorithm has a computational complexity that is linear in the size of the vector to be processed. This property makes it suitable also for real-time applications.

1 INTRODUCTION

Atrial fibrillation (AF) is a cardiac arrhythmia characterized by disorganized atrial electrical activity causing loss of effective contraction. AF is the most common persistent cardiac arrhythmia and it is also the most common cause of arrhythmia-related hospitalizations (Feinberg et al., 1995; Go et al., 2001). It has an enormous social impact because of its very high incidence and its clinical consequences. Moreover, it is often difficult to diagnose and its management is not optimized. The incidence of AF increases with age and given the life expectancies increasing in both developed and developing countries, AF is projected to become an important burden on most health care systems (McBride et al., 2008).

AF is the most common arrhythmia in the western countries, it is responsible for 70 to 100 thousand strokes per year in the US and is independently associated with up to 1.9-fold increase in the risk of death (Benjamin et al., 1998).

AF is also associated with extensive atrial structural, contractile, and electrophysiological remodeling, which can sustain AF itself (Nattel et al., 2008). Current pharmacological treatments of AF present some limits because they can be ventricular proarrhythmic and not able to prevent recurrences of AF.

A great deal of research on AF has focused on the identification of factors that can predict its first occurrence or recurrence. This could help in defining the best treatment in individual patients.

Promising results have been obtained from ECG signal processing, particularly from the analysis of P-waves. Electrocardiographic characteristics of AF have proven to be helpful in identifying patients at risk (Dilaveris and Gialafos, 2002; Platonov et al., 2000; Carlson et al., 2001; Bayes de Luna et al., 1999).

Prolonged and highly variable P-waves have been observed in patients prone to AF. Time-domain and morphological characteristics of P-waves from surface ECG recordings turned out to significantly distinguish patients at risk of AF (Bayes de Luna et al., 1999; Perez et al., 2009).

Most of these studies rely on measurements based on visual inspection; however, computerized automatic analyses can now be performed and have been recently introduced to estimate P-wave duration and morphological indices (Censi et al., 2007; Censi et al., 2008).

Due to the low signal-to-noise ratio associated with P-waves, this portion of ECG signal is usually analyzed performing signal averaging in order to build a P-wave template. Then, waves duration

and morphological features are extracted from this template. The unavoidable drawback of this approach is that some information is lost in the averaging operation. Each P-wave provides important information about the corresponding depolarization pattern throughout the atrial substrate. Thus, tracking changes between consecutive P-waves turns out to be extremely important in improving the understanding of the pathophysiological mechanisms of atrial substrates predisposing to AF.

However, the analysis of the temporal variability of consecutive P-waves is possible only if reliable beat-to-beat P-waves are available. This is attainable only if noise and artifacts are effectively filtered out from each single P-wave.

The aim of this investigation is to propose a novel method to smooth and denoise P-waves extracted from high-resolution DC-coupled ECG recordings, in order to improve the signal-to-noise ratio (SNR) of each single P-wave.

2 RATIONALE

From the previous section, it is evident that the ability to conduct a meaningful analysis of predisposing factors to AF strongly depends on the availability of reliable P-waves. In this regard, P-waves are reliable if the detrimental effects of noise and artifacts are reduced to an acceptable level.

In this section we propose an algorithm that is particularly effective for P-waves smoothing and denoising. It is based on the following idea. The measured P-wave is affected by noise and artifacts whose effect is to introduce additional “variability” into the observed P-wave with respect to the true one. Thus, provided that we introduce a suitable index of variability, smoothing and denoising can be performed by searching for a version of the P-wave that is close, in some sense, to the observed one, but has less “variability”. We make this idea precise in the following.

Denote by $\mathbf{z} = [z_1 \cdots z_n]^T$ the vector collecting n samples of a noisy P-wave extracted from an ECG tracing. The variability of a generic vector can be quantified introducing the following

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

$$[\mathbf{x}] \doteq \sum_{k=1}^{n-1} (x_k - x_{k+1})^2 \quad (1)$$

and is denoted by $[\mathbf{x}]$.

The quadratic variation is a well-known property used in the analysis of stochastic processes (Shreve,

2004). However, in this context we consider it as a function of deterministic or random vectors.

Introducing the $(n-1) \times n$ matrix

$$\mathbf{D} = \begin{bmatrix} 1 & -1 & 0 & \dots & 0 \\ 0 & 1 & -1 & \ddots & \vdots \\ \vdots & \ddots & \ddots & \ddots & 0 \\ 0 & \dots & 0 & 1 & -1 \end{bmatrix}, \quad (2)$$

the quadratic variation of \mathbf{x} can be expressed as

$$[\mathbf{x}] = \|\mathbf{D}\mathbf{x}\|^2, \quad (3)$$

where $\|\cdot\|$ denotes the Euclidean norm.

The quadratic variation is a consistent index of variability and its use is motivated by the following property. For vectors affected by additive noise, on average it does not decrease and moreover it is an increasing function of noise variances. In fact, let $\mathbf{x} = \mathbf{x}_0 + \mathbf{w}$, where \mathbf{x}_0 is a deterministic vector and $\mathbf{w} = [w_1 \cdots w_n]$ is a zero-mean random vector with covariance matrix $\mathbf{K}_w = \mathbb{E}\{\mathbf{w}\mathbf{w}^T\}$. We do not make any assumption about the distribution of \mathbf{w} , so the following considerations hold regardless of it. Computing the averaged quadratic variation of \mathbf{x} we get

$$\begin{aligned} \mathbb{E}\{\|\mathbf{D}\mathbf{x}\|^2\} &= \|\mathbf{D}\mathbf{x}_0\|^2 + \mathbb{E}\{\text{tr}(\mathbf{D}\mathbf{w}\mathbf{w}^T\mathbf{D}^T)\} \\ &= \|\mathbf{D}\mathbf{x}_0\|^2 + \text{tr}(\mathbf{D}\mathbf{K}_w\mathbf{D}^T) \end{aligned} \quad (4)$$

where, in the first equality, we have exploited the invariance of the trace under cyclic permutations. Note that $\text{tr}(\mathbf{D}\mathbf{K}_w\mathbf{D}^T) \geq 0$, since it is the trace of a positive semidefinite matrix (Horn and Johnson, 1990), but in all practical cases the inequality is strict. In fact, we have

$$\begin{aligned} \text{tr}(\mathbf{D}\mathbf{K}_w\mathbf{D}^T) &= \sum_{k=1}^{n-1} \mathbb{E}\{(w_k - w_{k+1})^2\} \\ &= \sum_{k=1}^{n-1} (\sigma_k^2 + \sigma_{k+1}^2 - 2\sigma_{k,k+1}) \end{aligned} \quad (5)$$

where $\sigma_k^2 = \mathbb{E}\{w_k^2\}$ and $\sigma_{k,k+1} = \mathbb{E}\{w_k w_{k+1}\}$. From (5) it follows that $\text{tr}(\mathbf{D}^T\mathbf{D}\mathbf{K}_w) = 0$ if and only if all the components of the noise vector \mathbf{w} are almost surely equal¹ and that $\mathbb{E}\{\|\mathbf{D}\mathbf{x}\|^2\}$ is an increasing function of noise variances.

For example, in typical scenarios of ECG tracings $\mathbf{w} = \mathbf{m} + \mathbf{a}$, where \mathbf{m} is due to white Gaussian noise whereas \mathbf{a} is due to the residual 50Hz or 60Hz power-line noise. We may assume $\mathbf{m} \sim \mathcal{N}(\mathbf{0}, \sigma_m^2 \mathbf{I})$ and $\mathbf{a} = [a_1 \cdots a_n]$ vector of samples from a harmonic process, i.e., $a_k = A \cos[2\pi \frac{f_0}{F_c} k + \phi]$, with A and ϕ

¹That is $w_1 = w_2 = \dots = w_n$ with probability 1.

independent, ϕ uniformly distributed in $[0, 2\pi)$, $f_0 \in \{50\text{Hz}, 60\text{Hz}\}$ and F_c being the sampling frequency. Moreover \mathbf{m} and \mathbf{a} are independent. In this case it is easy to verify that

$$\begin{aligned} \text{tr}(\mathbf{D}\mathbf{K}_w\mathbf{D}^T) &= \\ &= 2(n-1) \left[\sigma_m^2 + 2\mathbb{E}\{A^2\} \sin^2\left(\pi\frac{f_0}{F_c}\right) \right] \\ &= \frac{2\|\mathbf{x}_0\|^2(n-1)}{n} \left[\frac{1}{\text{SNR}} + \frac{4}{\text{SIR}} \sin^2\left(\pi\frac{f_0}{F_c}\right) \right] \quad (6) \end{aligned}$$

where $\text{SNR} = \frac{\|\mathbf{x}_0\|^2}{n\sigma_m^2}$ denotes the signal-to-noise ratio and $\text{SIR} = \frac{2\|\mathbf{x}_0\|^2}{n\mathbb{E}\{A^2\}}$ is the signal-to-interference ratio, considering the power-line noise as interference. From (6) it is evident that the average quadratic variation is a decreasing function of SNR and SIR.

In the following section we devise an efficient smoothing algorithm for P-waves exploiting the concept of quadratic variation.

3 SMOOTHING P-WAVES

In this section we denote by \mathbf{p} the vector collecting samples from the measured P-wave, the one that is affected by noise and artifacts, and by \mathbf{x} the corresponding vector after smoothing. Following the line of reasoning presented in the previous section, we determine \mathbf{x} solving the following optimization problem

$$\begin{cases} \text{minimize} & \|\mathbf{x} - \mathbf{p}\|^2 \\ \text{subject to} & \|\mathbf{D}\mathbf{x}\|^2 \leq a \end{cases} \quad (7)$$

where \mathbf{D} is defined in (2) and a is a positive constant that controls the degree of smoothness for \mathbf{p} . Its value is chosen in accordance with the peculiarity of the problem and satisfies $a < \|\mathbf{D}\mathbf{p}\|^2$ in order to avoid trivial solutions.² Note that we do not need to know in advance the appropriate value for a in any particular problem. In fact, as it will be clear later, the solution to the optimization problem (7) can be expressed in terms of a parameter that controls the degree of smoothness, i.e., the quadratic variation of the solution, and that is related to the value of a in (7). In this way, smoothing can be performed without caring about a in the constraint $\|\mathbf{D}\mathbf{x}\|^2 \leq a$, and reducing parametrically the quadratic variation of the solution to the desired level. In general, the optimal value for the controlling parameter can be found, as the one that entails the maximum SNR gain.

²When $a \geq \|\mathbf{D}\mathbf{p}\|^2$ the solution is $\mathbf{x} = \mathbf{p}$ and no smoothing is performed.

Let us consider (7) in more detail. It is a convex optimization problem, since both the objective function and the inequality constraint are convex. As a consequence, any locally optimal point is also globally optimal and Karush-Kuhn-Tucker (KKT) conditions provide necessary and sufficient conditions for optimality (Boyd and Vandenberghe, 2004). Moreover, since the objective function is strictly convex and the problem is feasible the solution exists and is unique. The Lagrangian is

$$\mathcal{L}(\mathbf{x}, \lambda) = \|\mathbf{x} - \mathbf{p}\|^2 + \lambda (\|\mathbf{D}\mathbf{x}\|^2 - a) \quad (8)$$

from the KKT conditions we get

$$\nabla \mathcal{L}(\mathbf{x}, \lambda) = 2(\mathbf{x} - \mathbf{p}) + 2\lambda \mathbf{D}^T \mathbf{D} \mathbf{x} = 0, \quad (9)$$

$$\lambda (\|\mathbf{D}\mathbf{x}\|^2 - a) = 0, \quad (10)$$

and the nonnegativity of the Lagrange multiplier $\lambda \geq 0$. However, if $\lambda = 0$ from (9) it follows that $\mathbf{x} = \mathbf{p}$, which is infeasible since the inequality constraint is not satisfied. Hence $\lambda > 0$ and from (10) it results that the inequality constraint is active, i.e., $\|\mathbf{D}\mathbf{x}\|^2 = a$. Thus, solving (9) we get eventually

$$\mathbf{x} = (\mathbf{I} + \lambda \mathbf{D}^T \mathbf{D})^{-1} \mathbf{p} \quad (11)$$

where \mathbf{I} denotes the identity matrix, and λ is determined by

$$\|\mathbf{D}\mathbf{x}\|^2 = \left\| \mathbf{D} (\mathbf{I} + \lambda \mathbf{D}^T \mathbf{D})^{-1} \mathbf{p} \right\|^2 = a. \quad (12)$$

Note that in (11) the inverse exists for any $\lambda \geq 0$ and when $\lambda = 0$ the solution corresponds to no smoothing. It is interesting that the solution to (11) is a linear operator acting on \mathbf{p} . Moreover, the Lagrange multiplier λ plays the role of a parameter controlling the degree of smoothing applied to \mathbf{p} . In fact, it is possible to prove that the quadratic variation of the solution, namely $\|\mathbf{x}\|^2 = \|\mathbf{D}\mathbf{x}\|^2$, is a *continuous and strictly decreasing* function of λ for $\lambda \in [0, +\infty)$ regardless of \mathbf{p} , provided that it is not a constant vector.³ This is equivalent to say that (12), when \mathbf{p} is not a constant vector, establishes a one-to-one correspondence between $a \in (0, \|\mathbf{D}\mathbf{p}\|^2]$ and $\lambda \in [0, +\infty)$, with $\lambda = 0$ corresponding to $a = \|\mathbf{D}\mathbf{p}\|^2$ and

$$\lim_{\lambda \rightarrow +\infty} \|\mathbf{x}\|^2 = 0. \quad (13)$$

A consequence of this is that we do not need to know in advance the value of a in (7), as smoothing can be performed according to (11) and λ can be

³If \mathbf{p} is a constant vector $\|\mathbf{x}\|^2 = \|\mathbf{p}\|^2 = 0$ regardless of λ .

adapted to fulfill some performance criterion. For example, considering the SNR gain⁴ as a quality index, λ can be chosen as the one that entails the maximum gain.

It is important to consider the computational aspects related to the smoothing operation, since a matrix inversion is involved in (11). If the size of the vector \mathbf{p} is large enough computational problems may arise. Actually this is not an issue for the typical length of vectors representing P-waves. However, if the same algorithm is applied to a complete ECG recording the computational burden, in terms of time and memory, and the accuracy become serious issues, even for batch processing.

It is possible to prove that due to the special structure of the matrix $\mathbf{I} + \lambda \mathbf{D}^T \mathbf{D}$, which is tridiagonal, smoothing in (11) can be performed with complexity $O(n)$, i.e., *linear in the size of vector \mathbf{p}* (Golub and Van Loan, 1996). This property is very important and makes the proposed algorithm suitable also for real-time applications.

Eventually, it is worthwhile noting that the algorithm we propose is not limited to P-wave smoothing, but it can be applied in very general situations, whenever smoothing and/or denoising are needed. This is due to the fact that the formulation and the rationale behind it, i.e., quadratic variation reduction, have general validity. In this regard, we successfully applied it also to EEG tracings denoising.

4 SIMULATION RESULTS

In order to evaluate performance of the proposed algorithm, a noiseless reference model of P-wave is needed. In this regard, we considered the P-wave model reported in (Censi et al., 2007), obtained fitting linear combinations of Gaussian functions to measured P-waves. Such a model can be reliably used to represent real P-waves as documented in the cited reference.

The noiseless reference model considered is reported in Figure 2. It has a duration of 200ms and its bandwidth essentially does not exceed 100Hz. It has been sampled at 2048Hz and the corresponding samples have been collected in a vector denoted by \mathbf{p}_0 .

The noiseless P-wave reference model \mathbf{p}_0 has been corrupted by additive noise, denoted by \mathbf{w} , where the components of \mathbf{w} are i.i.d.⁵ zero mean Gaussian ran-

⁴This is the ratio between the SNR after and before smoothing.

⁵Independent identically distributed.

dom variables with variance σ_w^2 . Thus, the corresponding noisy P-wave is

$$\mathbf{p} = \mathbf{p}_0 + \mathbf{w}. \quad (14)$$

In order to assess performance of the proposed algorithm the following quantities have been considered:

- the signal-to-noise ratio before smoothing

$$\text{SNR} = \frac{\|\mathbf{p}_0\|^2}{n \cdot \sigma_w^2} \quad (15)$$

where n is the size of vector \mathbf{p}_0 ;

- the signal-to-noise ratio after smoothing

$$\text{SNR}_s = \frac{\|\mathbf{p}_0\|^2}{\|\mathbf{x} - \mathbf{p}_0\|^2} \quad (16)$$

where $(\mathbf{x} - \mathbf{p}_0)$ is the error vector with respect to the reference model \mathbf{p}_0 .

In definition (16), we consider as noise affecting the smoothed vector \mathbf{x} , both the residual Gaussian noise and the reconstruction error.

Performances are measured in terms of SNR gain, defined as

$$G_{\text{snr}} = \frac{\text{SNR}_s}{\text{SNR}} = \frac{n \cdot \sigma_w^2}{\|\mathbf{x} - \mathbf{p}_0\|^2}. \quad (17)$$

Simulations have been carried out applying smoothing in accordance with (11) and choosing for λ the value, denoted by λ_{opt} , that entails the maximum SNR gain.

In Figure 2 the noiseless P-wave reference model \mathbf{p}_0 (dashed-line) is reported. Figure 1 shows the corresponding noisy version \mathbf{p} , corrupted by additive noise according to (14), with $\text{SNR} = 0\text{dB}$. Note that the SNR in this case is quite low, nevertheless the proposed algorithm is very effective in denoising \mathbf{p} and the resulting smoothed vector, namely \mathbf{x} , is plotted in Figure 2 (continuous-line) for an easy comparison with the reference model. In this case the resulting SNR gain is quite remarkable and amounts to $G_{\text{snr}} = 19.4\text{dB}$. It is important to point out that we evaluated the proposed algorithm on different models of P-wave and the resulting gains were all consistent with the ones reported in this work.

In order to evaluate how gain varies as input SNR changes, in Figure 3 we report the average SNR gain versus input SNR (bottom axis), when the reference P-wave model of Figure 2 is corrupted by additive noise with SNR ranging from -20dB to 40dB . The top axis of Figure 3 represents the corresponding in-band SNR, computed using the 100Hz bandwidth of the P-wave reference model. The in-band SNR is

about⁶ 10.1 dB greater than the corresponding input SNR.

For each input SNR we considered the solutions with $\lambda = \lambda_{opt}$ for 300 noise realizations and then we averaged the corresponding SNR gains. As Figure 3 highlights, the proposed algorithm exhibits a remarkable ability in smoothing P-waves. It achieves considerable gains over the whole range of practical input SNRs and for input SNR ≤ 6 dB the average gain exceeds 15 dB. It is remarkable that even when the SNR is quite high the algorithm exhibits considerable gain.

It is worth noting that the proposed algorithm is able to reject both out-of-band and in-band noise. In this regard, low-pass filtering cannot reject in-band noise without altering the signal. Indeed, an ideal 100Hz low-pass filter, in the same setting of our simulation, would exhibit a constant average gain of about 10.1 dB over the whole range of input SNR, as a result of rejection of the sole out-of-band noise. This is confirmed by simulation where we considered a linear-phase FIR low-pass filter synthesized applying the window method (Oppenheim et al., 1999) to an ideal 100Hz low-pass filter, using a Kaiser window and requiring 0.1 dB ripple in passband and 80 dB attenuation in stopband. In Figure 3 we report the filter SNR gain versus input SNR, averaged over the same 300 noise realizations used before.

Figure 3 highlights the effectiveness of the proposed algorithm and shows, in particular considering the in-band SNR, that it outperforms low-pass filtering for all practical values of SNR.




Figure 1: P-wave reference model of Figure 2 corrupted by additive Gaussian noise with SNR = 0 dB.




Figure 2: P-wave reference model p_0 (dashed-line) and smoothed solution x (continuous-line) from the noisy P-wave of Figure 1.

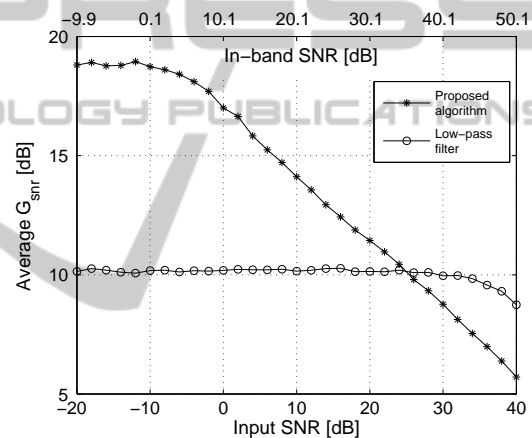


Figure 3: Average SNR gain G_{SNR} versus input SNR (bottom axis) and in-band SNR (top axis).

5 CONCLUSIONS

This work has been motivated by the need to have reliable beat-to-beat P-waves, as tracking changes between consecutive waves turns out to be very important in understanding the mechanisms underlying AF. This is attainable only if noise and artifacts are filtered out effectively from each single P-wave.

To solve this problem we have developed a smoothing and denoising algorithm. It is based on the notion of *quadratic variation* meant as a suitable index of variability for vectors or sampled functions.

The algorithm is the closed-form solution to a constrained convex optimization problem, where smoothing and denoising are achieved by reducing the quadratic variation of noisy P-waves. The computational complexity of the algorithm is *linear* in the

⁶Actually it is $10 \log(10.24)$ dB ≈ 10.1 dB, where 10.24 is the ratio between half of the sampling frequency and bandwidth of the P-wave reference model.

size of the vector to be processed, and this makes it suitable also for real-time applications. Simulation results confirm the effectiveness of the approach and highlight a remarkable ability to smooth and denoise P-waves.

Eventually, it is worthwhile noting that the proposed algorithm can be effectively applied to a wider range of signals, e.g., whole ECG or EEG tracings, whenever smoothing and/or denoising are needed.

REFERENCES

- Bayes de Luna, A., Guindo, J., Vinolas, X., Martinez-Rubio, A., Oter, R., and Bayes-Genis, A. (1999). Third-degree inter-atrial block and supraventricular tachyarrhythmias. *Europace*, 1:43–46.
- Benjamin, E. J., Wolf, P. A., D’Agostino, R. B., Silbershatz, H., Kannel, W. B., and Levy, D. (1998). Impact of atrial fibrillation on the risk of death: the Framingham Heart Study. *Circulation*, 98:946–952.
- Boyd, S. and Vandenberghe, L. (2004). *Convex Optimization*. Cambridge University Press.
- Carlson, J., Johansson, R., and Olsson, S. B. (2001). Classification of electrocardiographic P-wave morphology. *IEEE Trans Biomed Eng*, 48:401–405.
- Censi, F., Calcagnini, G., Ricci, C., Ricci, R. P., Santini, M., Grammatico, A., and Bartolini, P. (2007). P-wave morphology assessment by a gaussian functions-based model in atrial fibrillation patients. *IEEE Trans Biomed Eng*, 54:663–672.
- Censi, F., Ricci, C., Calcagnini, G., Triventi, M., Ricci, R. P., Santini, M., and Bartolini, P. (2008). Time-domain and morphological analysis of the P-wave. Part I: Technical aspects for automatic quantification of P-wave features. *Pacing Clin Electrophysiol*, 31:874–883.
- Clifford, G. D., Azuaje, F., and McSharry, P. (2006). *Advanced Methods And Tools for ECG Data Analysis*. Artech House, Inc., Norwood, MA, USA.
- Dilaveris, P. E. and Gialafos, J. E. (2002). Future concepts in P wave morphological analyses. *Card Electrophysiol Rev*, 6:221–224.
- Feinberg, W. M., Blackshear, J. L., Laupacis, A., Kronmal, R., and Hart, R. G. (1995). Prevalence, age distribution, and gender of patients with atrial fibrillation. Analysis and implications. *Arch. Intern. Med.*, 155:469–473.
- Go, A. S., Hylek, E. M., Phillips, K. A., Chang, Y., Henault, L. E., Selby, J. V., and Singer, D. E. (2001). Prevalence of diagnosed atrial fibrillation in adults: national implications for rhythm management and stroke prevention: the AnTicoagulation and Risk Factors in Atrial Fibrillation (ATRIA) Study. *JAMA*, 285:2370–2375.
- Golub, G. H. and Van Loan, C. F. (1996). *Matrix computations*. Johns Hopkins University Press, Baltimore, MD, USA, 3 edition.
- Horn, R. A. and Johnson, C. R. (1990). *Matrix Analysis*. Cambridge University Press.
- McBride, D., Mattenklotz, A. M., Willich, S. N., and Bruggenjurg, B. (2008). The Costs of Care in Atrial Fibrillation and the Effect of Treatment Modalities in Germany. *Value Health*.
- Nattel, S., Burstein, B., and Dobrev, D. (2008). Atrial remodeling and atrial fibrillation: mechanisms and implications. *Circ Arrhythm Electrophysiol*, 1(1):62–73.
- Oppenheim, A. V., Schafer, R. W., and Buck, J. R. (1999). *Discrete-time signal processing (2nd ed.)*. Prentice-Hall, Inc.
- Perez, M. V., Dewey, F. E., Marcus, R., Ashley, E. A., Al-Ahmad, A. A., Wang, P. J., and Froelicher, V. F. (2009). Electrocardiographic predictors of atrial fibrillation. *Am. Heart J.*, 158:622–628.
- Platonov, P. G., Carlson, J., Ingemansson, M. P., Roijer, A., Hansson, A., Chireikin, L. V., and Olsson, S. B. (2000). Detection of inter-atrial conduction defects with unfiltered signal-averaged P-wave ECG in patients with lone atrial fibrillation. *Europace*, 2:32–41.
- Shreve, S. E. (2004). *Stochastic Calculus for Finance II: Continuous-Time Models*. Springer Finance. Springer Science+Business Media, Inc.